J. Eur. Math. Soc. (Online first) DOI 10.4171/JEMS/1218

Jonathan Hanselman

Heegaard Floer homology and cosmetic surgeries in S^3

Received December 13, 2019

Abstract. If a knot K in S^3 admits a pair of truly cosmetic surgeries, we show that the surgery slopes are either ± 2 or $\pm 1/q$ for some value of q that is explicitly determined by the knot Floer homology of K. Moreover, in the former case the genus of K must be 2, and in the latter case there is a bound relating q to the genus and the Heegaard Floer thickness of K. As a consequence, we show that the cosmetic crossing conjecture holds for alternating knots (or more generally, Heegaard Floer thin knots) with genus not equal to 2. We also show that the conjecture holds for any knot K for which each prime summand of K has at most 16 crossings; our techniques rule out cosmetic surgeries in this setting except for slopes ± 1 and ± 2 on a small number of knots, and these remaining examples can be checked by comparing hyperbolic invariants. These results make use of the surgery formula for Heegaard Floer homology, which has already proved to be a powerful tool for obstructing cosmetic surgeries; we get stronger obstructions than previously known by considering the full graded theory. We make use of a new graphical interpretation of knot Floer homology and the surgery formula in terms of immersed curves, which makes the grading information we need easier to access.

Keywords. Dehn surgery, cosmetic surgery, knot Floer homology

1. Introduction

Given a knot K in S^3, two surgeries $S^3_r(K)$ and $S^3_{r'}(K)$ with $r \neq r'$ are said to be *cosmetic* if $S_r(K)$ and $S_{r'}(K)$ are diffeomorphic as unoriented manifolds, and *truly cosmetic* if $S_r(K) \cong S_{r'}(K)$ (here, and throughout the paper, $\cong$ denotes orientation preserving diffeomorphism). Surgeries that are cosmetic but not truly cosmetic are called *chirally cosmetic*. Cosmetic surgeries are one way in which the surgery characterization of a 3-manifold can fail to be unique. Examples of chirally cosmetic surgeries are not difficult to find, but Gordon conjectured that there are no truly cosmetic surgeries on nontrivial knots [4, Conjecture 6.1] (see also [14, Problem 1.81 A]). This conjecture is stated more generally for knots in arbitrary 3-manifolds, with the notion of truly cosmetic surgery suitably extended, but we will only consider the case of knots in S^3.

Jonathan Hanselman: Department of Mathematics, Princeton University, Princeton, NJ, USA; jh66@princeton.edu

Mathematics Subject Classification (2020): 57K18, 57K31

Conjecture 1 (Cosmetic Surgery Conjecture in S^3). *Let K be a nontrivial knot in S^3. If $r \neq r'$, then $S^3_r(K) \ncong S^3_{r'}(K)$.*

The conjecture can be viewed as a generalization of the knot complement problem, solved by Gordon and Luecke [5], who proved that no pair of cosmetic surgeries contains the trivial surgery $S^3_\infty(K)$. In addition to this, several partial results related to Conjecture 1 are known. Boyer and Lines [1] used surgery formulae for Casson–Walker and Casson–Gordon invariants to place a restriction on $\Delta_K(t)$ for knots K admitting truly cosmetic surgeries. Much of the recent progress on Conjecture 1 has made use of Heegaard Floer homology, which has led to several results on obstructing cosmetic or truly cosmetic surgeries. For any pair of truly cosmetic surgeries, the surgery slopes were shown first to have opposite signs [19, 24], and then to be in fact opposite [15]. If K admits truly cosmetic surgeries, then the genus of K is not 1 [23] and the knot Floer homology of K satisfies certain additional constraints [3, 15].

Heegaard Floer homology has already proved to be a powerful tool at distinguishing surgeries, but it has not been used to its full potential. Each application to the cosmetic surgery conjecture mentioned above uses only partial information from Heegaard Floer homology, either the total rank of $\widehat{HF}$ [19,24], the d-invariants [15,23], or the Euler characteristic of HF^+_{red} [3, 15]. We will harness (almost) all of the information in Heegaard Floer homology to obtain much stronger obstructions to truly cosmetic surgeries. In particular, we consider the isomorphism type of $\widehat{HF}$ as an absolutely graded vector space, which amounts to keeping track of the grading for each generator in addition to the rank. This is facilitated by a recent reinterpretation of Heegaard Floer invariants for manifolds with torus boundary in terms of collections of immersed curves due to the author, Rasmussen, and Watson [8, 9]. In particular, this provides a combinatorial framework that makes comparing gradings for surgeries on knots easier. We give obstructions to truly cosmetic surgeries in terms of numerical invariants n_s extracted from knot Floer homology, the Heegaard Floer thickness $\mathrm{th}(K)$, and the Seifert genus $g(K)$. The Heegaard Floer thickness is the difference between the maximal and minimal δ-gradings in $\widehat{HFK}$. The integers n_s will be introduced in Section 2.3; briefly, they count the intersection number of the immersed multicurve representing knot Floer homology with a horizontal line at height i. Our main result is the following:

Theorem 2. *If $K \subset S^3$ is a nontrivial knot and $S^3_r(K) \cong S^3_{r'}(K)$ for $r \neq r'$, then*

(i) *the pair of slopes $\{r, r'\}$ is either $\{\pm 2\}$ or $\{\pm 1/q\}$ where $q = \dfrac{n_0 + 2\sum_{s=1}^{\infty} n_s}{4\sum_{s=1}^{\infty} s^2 n_s}$;*

(ii) *if $\{r, r'\}$ is $\{\pm 2\}$ then $g(K) = 2$ and $n_0 = 2n_1$;*

(iii) *if $\{r, r'\}$ is $\{\pm 1/q\}$ then $q \leq \dfrac{\mathrm{th}(K) + 2g(K)}{2g(K)(g(K)-1)}$.*

Note that for any given knot we rule out all but at most two pairs of truly cosmetic surgeries; it was not previously known that a knot must have finitely many pairs of truly cosmetic surgeries. Importantly, this is an *effective* finiteness statement, meaning that for a given knot K the two (or fewer) potential truly cosmetic surgery pairs are explicitly

determined. This makes it possible to check the conjecture for a given finite set of knots by means of a finite computation (see Theorem 6 below). In fact, in the vast majority of cases observed, the value of q predicted by conclusion (i) is not an integer and the conditions in (ii) are not met, so that Theorem 2 rules out all truly cosmetic surgeries on K.

Although the formula in conclusion (i) of Theorem 2 determines q, the bound on q in conclusion (iii) is also useful as a convenient way to rule out truly cosmetic surgeries on large classes of knots without computing the n_s invariants. In particular:

Corollary 3. *If K is a nontrivial knot in S^3 with $g(K) \neq 2$ and $\mathrm{th}(K) < 6$, then K does not admit any truly cosmetic surgeries.*

Proof. The cosmetic surgery conjecture is known for genus 1 knots [23], so we may assume that $g(K) > 2$. Conclusion (ii) from Theorem 2 rules out slopes ± 2. The bound from conclusion (iii) of Theorem 2 then implies that $q < 1$, which is impossible. ∎

A great many knots have $\mathrm{th}(K) < 6$. A knot is *Floer homologically thin* if $\mathrm{th}(K) = 0$, that is, if only one δ-grading is occupied (we will simply refer to such knots as *thin*); examples of thin knots include all alternating and quasialternating knots. Furthermore, direct computation reveals that for any prime knot K with at most 16 crossings, $\mathrm{th}(K) \leq 2$. It follows that the cosmetic surgery conjecture holds for any of these knots with genus other than 2.

An additional consequence of Theorem 2 is a restriction on the manifolds that could arise from truly cosmetic surgeries on knots in S^3:

Corollary 4. *If Y is a closed oriented 3-manifold with $|H_1(Y;\mathbb{Z})| > 2$, then Y cannot be obtained by a truly cosmetic surgery on any knot in S^3.*

In addition to those in Theorem 2, some further conditions must also be satisfied by the knot Floer homology of K for truly cosmetic surgeries to exist. These conditions require some more definitions to state in general (see Section 5), but in the case of thin knots they can be stated in terms of the Alexander polynomial and signature.

Theorem 5. *If a nontrivial knot $K \subset S^3$ is thin (in particular if K is alternating or quasialternating) and admits a pair of truly cosmetic surgeries, then $\Delta_K(t) = nt^2 - 4nt + (6n+1) - 4nt^{-1} + nt^{-2}$ for some positive integer n, $\sigma(K) = 0$, and the surgery slopes are ± 1 or ± 2.*

We remark that Theorem 5 is the best statement possible for thin knots using only Heegaard Floer homology. That is, if K is thin, $\Delta_K(t) = nt^2 - 4nt + (6n+1) - 4nt^{-1} + nt^{-2}$, and $\sigma(K) = 0$, then the pairs $\{S^3_1(K), S^3_{-1}(K)\}$ and $\{S^3_2(K), S^3_{-2}(K)\}$ are not distinguished by their Heegaard Floer homology. The first two examples in the knot tables where this occurs are the knots 9_{41} and 9_{44}. Similarly, Theorem 2 and the additional constraints in Section 5 allow us to extract as much information as possible from Heegaard Floer homology: if we cannot prove the cosmetic surgery conjecture for a given knot,

then the pairs of surgeries that are not ruled out in fact have isomorphic Heegaard Floer homology.

Knots having any surgeries at all that cannot be distinguished by Heegaard Floer homology are exceedingly rare, but they do exist—to date we have found 337 such knots, each with two pairs of slopes that are not distinguished. All of these 337 examples are genus 2 and have Alexander polynomial of the form described in Theorem 5 (though not all are thin). To prove the conjecture for these knots, we must use other invariants to distinguish the remaining pairs of surgeries; for the examples found so far the hyperbolic volume and Chern–Simons invariant are sufficient. In this way we verify the conjecture for all prime knots up to 16 crossings, and in fact for arbitrary connected sums of such knots.

Theorem 6. *Let $K \subset S^3$ be a nontrivial knot whose prime summands each have at most 16 crossings. If $r \neq r'$ then $S^3_r(K) \ncong S^3_{r'}(K)$.*

This paper grew out of an attempt to answer the question: how much can Heegaard Floer homology tell us about the Cosmetic Surgery Conjecture? For knots in S^3, we have now given a comprehensive answer to that question. Indeed, we see that Heegaard Floer homology can say a great deal, and examples for which it is not sufficient to prohibit truly cosmetic surgeries appear to be very rare. Nevertheless, other tools will be required to prove Conjecture 1 outright.

The results in this paper are, to the author's knowledge, the strongest obstructions available for knots in S^3, but similar results can be obtained using other techniques. In particular, very recently Futer, Purcell, and Schleimer [2] have used hyperbolic methods to prove a result comparable to Theorem 2: for any given hyperbolic knot, they rule out truly cosmetic surgeries on all but an explicitly determined finite set of slopes. For knots in S^3 this result seems to be weaker in practice than Theorem 2, in the sense that the finite set of slopes remaining is larger and thus the exhaustive search required to check the conjecture on a given set of knots is slower. However, the result in [2] applies not just to hyperbolic knot complements in S^3 but to arbitrary finite volume hyperbolic manifolds with torus boundary. In contrast, the arguments in this paper are highly specialized to knots in S^3 (though S^3 can be replaced with any integer homology sphere L-space). It is interesting to ask how much Heegaard Floer homology can tell us about cosmetic surgeries in more general manifolds. Although the proofs in this paper are not well suited to that setting, the broader principle of using immersed curves to more easily compare the ranks and relative gradings of the Heegaard Floer invariants of different Dehn fillings may be fruitful. In particular, we could hope to obtain a finiteness result in the line of Theorem 2 and [2, Theorem 7.29].

Question 7. Can Heegaard Floer homology be used to rule out all but finitely many pairs of cosmetic surgery slopes for arbitrary manifolds with torus boundary?

A good starting point would be graph manifolds with torus boundary, since (i) hyperbolic techniques would not apply and (ii) there is well developed machinery for understanding the bordered Floer invariants in this case [6, 10].

The rest of the paper is organized as follows. In Section 2 we describe knot Floer homology and review its relevant properties. This section is recommended even for readers already familiar with knot Floer homology, as our description of the invariant is not the usual one. In particular, we describe knot Floer homology as a decorated collection of immersed curves, a perspective that we will use throughout the paper. Section 3 addresses the Cosmetic Surgery Conjecture and briefly reviews some existing results; this is not meant to be a comprehensive survey of the subject, but rather focuses on results that use Heegaard Floer homology and on which our arguments build. In Section 4 we introduce our main obstructions and prove Theorem 2. Section 5 refines these results and provides several explicit obstructions to a knot admitting truly cosmetic surgeries; in particular, we prove Theorem 5. Finally, in Section 6 we verify the conjecture for arbitrary connected sums of knots up to 16 crossings, proving Theorem 6.

2. Knot Floer homology

Knot Floer homology was defined by Oszváth and Szabó [17] and independently by Rasmussen [21]. We will use a description of this invariant for knots in S^3 in terms of immersed curves; this is rather different from the original formulation, though it carries equivalent information. We will be primarily interested in a weaker form of the invariant, which we call $\widehat{\Gamma}(K)$ and which is equivalent to the $UV = 0$ truncation of the knot Floer complex. The $UV = 0$ truncation of knot Floer homology is also equivalent to bordered Floer homology of the knot complement, and an immersed curve description of this invariant is due to the author, Rasmussen, and Watson [8, 9] (the case of knot complements is discussed specifically in [9, Section 4]). In particular, the invariant denoted $\widehat{\Gamma}(K)$ in this paper agrees with $\widehat{HF}(M)$ with $M = S^3 \setminus \nu(K)$ in the notation of [8, 9]. For readers unfamiliar with bordered Floer homology, a bordered free construction of the immersed curves $\widehat{\Gamma}(K)$ will appear in a forthcoming paper by the author [7]. This construction has the advantage that it can be strengthened to a decorated curve $\Gamma(K)$ capturing the full knot Floer complex $CFK^\infty(K)$. We will not need this stronger invariant in the present paper, though we will need to make use of the construction in [7] in one small way (see Proposition 17).

We will now describe the invariant $\widehat{\Gamma}(K)$. Throughout we work with coefficients in $\mathbb{F} = \mathbb{Z}/2\mathbb{Z}$. We begin by setting notation for the spaces in which the curves $\widehat{\Gamma}(K)$ appear. Let T denote the torus marked with a chosen pair of parametrizing curves μ and λ and a single marked point w, which we may take to be the intersection of μ and λ. Let $\overline{T}$ denote the infinite cyclic covering space of T in which λ lifts to a loop and μ does not, and let $\bar{p} : \overline{T} \to T$ denote the covering map. We will identify $\overline{T}$ with $(\mathbb{R}/\mathbb{Z}) \times \mathbb{R}$, where the lifts of λ and μ are horizontal and vertical, respectively, and the preimages of w are the points $(0, s - 1/2)$ for integers s. Let $\widetilde{T}$ denote the universal covering space $\mathbb{R}^2$ with covering map $\widetilde{p} : \widetilde{T} \to \overline{T}$. By slight abuse, we will often refer to the vertical line through the marked points in $\overline{T}$ (or through a column of marked points in $\widetilde{T}$) as μ, though it is really a lift of the curve μ in T. Finally, we will use $T_\bullet$, $\overline{T}_\bullet$, and $\widetilde{T}_\bullet$ to denote corresponding punctured

surfaces obtained by removing the marked points. We may conflate punctures and marked points at times, the only distinction is that we use marked points if we want to allow disks to cover these points and we use punctures otherwise.

2.1. *The knot Floer invariant.*

To a knot K in S^3, we associate a collection $\boldsymbol{\gamma} = \{\gamma_0, \dots, \gamma_n\}$ of oriented immersed curves in $\overline{T}_\bullet$; $\boldsymbol{\gamma}$ is an invariant knot K, up to regular homotopy and reindexing of the curves (note that we work in the punctured cylinder $\overline{T}_\bullet$ rather than marked cylinder $\overline{T}$, meaning that the curves and homotopies are required to avoid the punctures). $\widehat{\Gamma}(K)$ denotes this multicurve along with two extra decorations: (1) each curve may be decorated with a local system, and (2) the multicurve carries decorations to encode relative grading data. We remark that the first decoration is not relevant to the arguments in this paper, though we describe it here for completeness. Moreover, it is still unknown whether the local system decoration is nontrivial for any knot in S^3. In contrast, the second type of decoration will play a crucial role. This Maslov grading decoration for immersed curves in the context of bordered Floer homology is discussed in detail in [9, Section 2]; see also [7] for the special case of $\widehat{\Gamma}(K)$.

Local systems. If a curve γ_i is homologous in $\overline{T}_\bullet$ to k_i copies of some primitive curve γ_i', then we will assume that γ_i is realized by k_i parallel copies of γ_i' outside of a small region, in which the curve crosses itself $k_i - 1$ times as shown in Figure 1 (a). This region contains one segment of γ_i (the negatively sloped segment in the figure) that intersects each of $k_i - 1$ other segments. Then each curve γ_i is decorated with a subset of these $k_i - 1$ self-intersection points. Note that for a primitive curve this decoration is automatically trivial. The selected intersection points should be interpreted as places where a traveler along the negatively sloped segment is allowed to make a left turn onto one of the other segments. In the language of [8], this means that we extend the curve γ_i to an immersed train track and add two (oriented) edges near the selected intersection point, as shown in Figure 1 (b). Decorating a chosen subsection of intersection points by adding train track edges in this way determines a $k_i \times k_i$ invertible matrix with coefficients in $\mathbb{F}$ that counts

Fig. 1. We assume that nonprimitive curves are arranged as shown in (a) and decorated with a subset of the intersection points in the boxed region. To each chosen intersection point we add a pair of edges as pictured in (b), or equivalently a crossover arrow in the notation of [8]. Counting (smooth) paths from the left side of the boxed region to the right determines a matrix, which can be interpreted as a local system.

immersed paths from the left to the right of the boxed region. By a local system we mean a similarity class of such matrices, which is equivalent to the subset of intersection points above since the matrix constructed in this way will be in rational canonical form. The local system associated to each curve in $\boldsymbol{\gamma}$ is also an invariant of K. Note that the pair of train track edges added at an intersection point is equivalent to a single "crossover arrow", in the shorthand notation of [8], and using the arrow sliding moves the whole configuration can be replaced with k_i parallel copies of γ_i' with some crossover arrows between parallel strands. This also defines a matrix, which is similar to the one constructed above.

Maslov gradings. The multicurve $\boldsymbol{\gamma}$ can be enhanced with extra decorations in order to encode some (relative) grading information. In fact, the desired information is already contained in the immersed curve for any single component γ_i of $\boldsymbol{\gamma}$, so extra decorations are only required to capture relative gradings between different components. There are multiple ways to encode this information the approach we describe here is to add labeled arcs to $\boldsymbol{\gamma}$ connecting different components. More precisely, we extend the multicurve $\boldsymbol{\gamma}$ to an immersed graph $\boldsymbol{\gamma}_{\mathrm{gr}}$, which contains $\boldsymbol{\gamma}$ as a subgraph and all of whose vertices are contained in $\boldsymbol{\gamma}$, but which also contains some number of edges connecting vertices on different curve components. We will refer to these new edges as *grading edges*, and they should be ignored except for the purposes of computing gradings. We require that the grading edges are tangent to $\boldsymbol{\gamma}$ at their endpoints, so that $\boldsymbol{\gamma}_{\mathrm{gr}}$ is in fact an immersed train track (recall that a *train track* is a graph for which all incident edges at any vertex are mutually tangent). Moreover, we require the ends of the grading edges to be consistent with the relative orientations on the curves, in the sense that a smooth path that runs over an edge connecting γ_i and γ_j either follows the orientation on both curves or opposes the orientation on both curves. Grading edges themselves are directed (this direction is not required to agree with the orientation on the curves) and labeled with an integer weight.

We say that a set of grading edges on $\boldsymbol{\gamma}$ is *complete* if $\boldsymbol{\gamma}_{\mathrm{gr}}$ is connected as a graph. We say that a set of grading edges is *consistent* if, for any closed (not necessarily smooth) path P in $\boldsymbol{\gamma}_{\mathrm{gr}}$,

$$-\operatorname{rotation}(P) + \operatorname{winding}(P) + \operatorname{weights}(P) = 0, \tag{1}$$

where $\operatorname{rotation}(P)$ is $\frac{1}{2\pi}$ times the total counterclockwise rotation along the smooth sections of P, $\operatorname{winding}(P)$ is the sum over marked points w in $\overline{T}$ of the winding number of P around w, and $\operatorname{weights}(P)$ is the sum of the weights of all grading edges traversed by P, with the weight counted negatively if P traverses the grading edge backwards. More precisely, for the rotation and winding numbers to make sense, we only consider paths P that do not wrap around the cylinder; these can be viewed as paths in the marked strip obtained by cutting $\overline{T}$ open along the line $\{1/2\} \times \mathbb{R}$. We say that two complete sets of grading edges are *equivalent* if their union is consistent. With these definitions established, the grading decoration we will use on the multicurve $\boldsymbol{\gamma}$ to define $\widehat{\Gamma}(K)$ is a complete consistent set of grading edges; this decoration is an invariant of K up to equivalence of sets of grading edges.

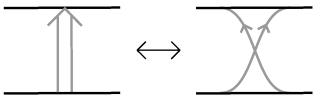

Fig. 2. A crossover arrow is shorthand for a pair of edges.

For any complete set of grading edges and any additional oriented grading edge, there is a unique integer weight on the new edge that makes the combined set consistent. In particular, for any oriented edge connecting two components of $\boldsymbol{\gamma}$, tangent to $\boldsymbol{\gamma}$ at each end in a way that is consistent with the orientation of the curves, we can choose a set of grading edges that contains this edge with some integer weight. If $\boldsymbol{\gamma}$ contains $n+1$ curves, then a minimal complete subset of any complete set of grading edges contains n edges. We will generally choose a minimal complete set of grading edges, except that the convention established in [9] replaces each grading edge with a pair of edges as in Figure 2, the shorthand for which is a bold arrow. We will call these pairs of grading edges *grading arrows*, and we will decorate $\boldsymbol{\gamma}$ with a minimal consistent set of grading arrows. Using both edges in a grading arrow is unnecessary but can be convenient. Note that a grading arrow can be labeled by a single weight, since consistency requires both edges in the pair to have the same weight. For examples of multicurves decorated with grading arrows, see Figure 3.

Remark 8. Grading arrows are described in detail in [9, Section 2] (see also [7] for a slightly different description). Note that in [9], grading arrows do not carry an integer weight; these arrows should be interpreted as having weight 0. By sliding arrows over punctures and changing the corresponding weights by 1, it is clear that any configuration of weighted arrows can be replaced with an equivalent configuration of arrows that all have weight 0. Thus using weighted grading arrows is not necessary, but it is convenient as it provides greater freedom in which arrows we choose. Finally, note that the arrow weights discussed here should not be confused with the complexity weights on crossover arrows used in the proof of the arrow removal algorithm in [8, Section 3.7].

2.2. *Properties of* $\widehat{\Gamma}$

Some examples of the invariant $\widehat{\Gamma}(K)$ are shown in Figure 3. For the unknot and the right hand trefoil, the invariant contains a single curve. The invariant for the figure eight knot consists of two curves decorated with a single grading arrow connecting them, while the invariant for 9_{44} has five curves and four grading arrows. These examples demonstrate some general properties of $\widehat{\Gamma}$, which we now highlight.

(I) For a knot K in S^3, the multicurve $\boldsymbol{\gamma}$ associated with $\widehat{\Gamma}$ can be homotoped to have exactly one intersection with the line $\{1/2\}\times\mathbb{R}$ in $\overline{T}\simeq(\mathbb{R}/\mathbb{Z})\times\mathbb{R}$. This follows from the fact that when $\widehat{\Gamma}$ and this vertical line are in minimal position, their intersections generate $\widehat{HF}$ of the meridional filling, that is, of $\widehat{HF}(S^3)\cong\mathbb{F}$ [8, Theorem 2]. This condition implies that there is one distinguished curve component, which we always take to be γ_0, that wraps around the cylinder once, and all other components can be contained in a neighborhood of μ, the vertical line through the marked points. As in the examples in Figure 3, γ_0 is always oriented left-to-right.

(II) The immersed curves in $\boldsymbol{\gamma}$ are *unobstructed*, meaning that they do not bound any "teardrops", or one-cusped disks, that do not enclose a marked point. That is, there are no immersed disks in the punctured cylinder $\overline{T}_\bullet$ with boundary on some $\gamma_i \in \boldsymbol{\gamma}$ such that the boundary is a smooth path apart from one acute corner at a self-intersection point of γ_i.

(III) The consistency condition for sets of grading arrows in (1) is stated as a condition that must hold for all closed loops in the train track $\boldsymbol{\gamma}_{\text{gr}}$ that do not wrap around the puncture. We remark that this condition must in particular hold for each curve component γ_i with $i \neq 0$, even before grading arrows are introduced, and this places restrictions on the allowed curves: for any closed curve with net zero rotation, the total winding number around punctures must also be zero. In particular, any figure eight shaped curve must enclose the same number of marked points on each side.

Remark 9. In the examples above, all of the curves γ_i with $i \neq 0$ are figure eights wrapping around two adjacent punctures; such a curve will be called a *simple figure eight*. This is not a general property of $\widehat{\Gamma}$, but it is incredibly common. In fact, this condition holds for all but one prime knot up to 15 crossings. The unique exception is $15n166130$, for which $\widehat{\Gamma}$ contains (along with thirty simple figure eights) two components not of this form. These components are still figure eight curves enclosing one marked point on each side, but they enclose nonadjacent marked points. Larger examples can be constructed with figure eight curves enclosing more than one marked point on each side, but the author has not yet found an example with a homologically trivial curve that is not a figure eight in this more general sense.

(IV) The decorated curve set $\widehat{\Gamma}(K)$ is invariant under rotation by π about the origin, up to homotopy of curves and equivalence of grading arrows, except that the rotation flips the orientation of every curve. This is the geometric expression of a symmetry for bordered Floer invariants established in [9, Theorem 7], which was proved earlier in the case of knot complements by Xiu [25].

2.3. Invariants derived from knot Floer homology

Several interesting numerical invariants of K can be extracted from $\widehat{\Gamma}$. For example, the genus of K is the maximum height of an intersection of $\widehat{\Gamma}$ with the vertical line μ through the marked points, assuming $\widehat{\Gamma}$ is in minimal position with μ. Here we mean height in the discrete sense: an intersection point is said to occur at height s if its y-coordinate falls between the marked points at $(0, s - 1/2)$ and $(0, s + 1/2)$.

There is a distinguished curve component γ_0 that wraps around the cylinder $\overline{T}$ and a distinguished intersection of γ_0 with μ, the first time γ_0 reaches μ after wrapping around the cylinder. The Ozsváth–Szabó τ invariant is the height of this first intersection point on γ_0 (see [9, Section 4.2]; this intersection corresponds to the generator of vertical homology and the height gives its Alexander grading). Moreover, after γ_0 reaches this first intersection point, it can do one of three things: turn right (downward), turn left (upward), or continue straight. Hom's invariant $\epsilon(K)$ is 1, -1, or 0, respectively, in these three cases (again see [9, Section 4.2]; this behavior corresponds to the generator of ver-

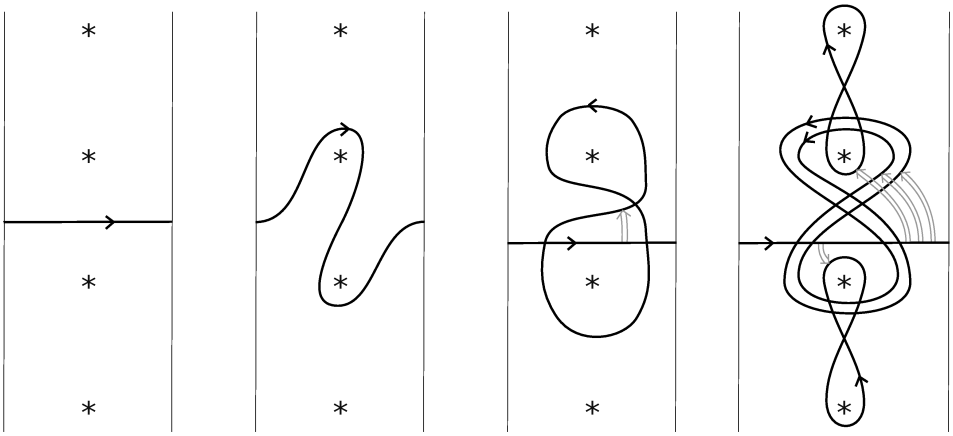

Fig. 3. The curve invariants $\widehat{\Gamma}(K)$ for, from left to right, the unknot, the right-handed trefoil, the figure eight knot, and 9_{44}. Grading arrows, shown in gray, all have weight 0, and there are no local systems.

tical homology lying at the end of a horizontal arrow, the beginning of a horizontal arrow, or no horizontal arrow). Note that by symmetry γ_0 can only continue straight if the intersection was at height 0, so $\epsilon(K) = 0$ implies $\tau(K) = 0$. In this case γ_0 intersects μ only once and is homotopic to the simple horizontal line $S^1 \times \{0\}$ in the cylinder. Note that τ and ϵ are both concordance invariants of K; in fact, it can be shown that the distinguished curve component γ_0 is itself a concordance invariant of K, up to homotopy of curves (this follows from [12, Theorem 1]; γ_0 is the immersed curve representative of the direct summand of $UV = 0$ knot Floer homology supporting the homology, which is a concordance invariant). In fact the curve γ_0 exactly encodes the ϵ-equivalence class as defined by Hom. Note that τ and ϵ depend only on γ_0.

In the arguments in this paper, it will be useful to quantify a few more aspects of the underlying multicurve $\boldsymbol{\gamma}$ for $\widehat{\Gamma}$. To do this, we will assume that $\boldsymbol{\gamma}$ has a convenient form. We start by "pulling tight" as described in [8]. This means that we assume $\boldsymbol{\gamma}$ is the minimal length representative of its homotopy class, subject to the constraint that it avoids an open disk of some small radius ϵ around each marked point. Intuitively, we think of there being a peg of radius ϵ at each marked point and think of the curve as a rope winding through the pegs; we pull the curve taut as if the rope were elastic.[1] Under this assumption, $\boldsymbol{\gamma}$ breaks into segments that connect pegs (separated by small portions of $\boldsymbol{\gamma}$ lying on the boundary of a peg). Because of property (I) from the previous section, exactly one of these segments leaves a neighborhood of μ and wraps around the cylinder, while all other segments connect a peg to one directly above it and thus are (roughly) vertical of length 1. We are interested in counting these vertical segments. We say that

[1] To ensure transverse self-intersection (and make pictures easier to read), we can modify this by letting the curve wrap at a slightly different radius each time it encounters a peg.

a vertical segment is at height s if it connects a peg at height $s - 1/2$ to a peg at height $s + 1/2$; let n_s denote the number of vertical segments at height s and let $n = \sum_{s\in\mathbb{Z}} n_s$ denote the total number of vertical segments. In general we will perturb the curve slightly from its pulled tight position so that the segments counted by the n_s are in fact vertical outside of a neighborhood of the pegs; Figure 4 shows a simple figure eight curve in this position. Note that a simple figure eight curve always contributes two vertical segments at the same height. The unique nonvertical segment in $\boldsymbol{\gamma}$ is also of interest, and we will record the slope m of this segment. For example, consider the curve invariant for the knot 9_{44} shown in Figure 3. Each figure eight curve contributes two vertical segments, and we have $m = 0$, $n_0 = 4$, $n_1 = n_{-1} = 2$, and $n_s = 0$ for all other s. Since they are derived from $\widehat{\Gamma}$, the quantities m and n_s are invariants of the knot, though we remark that they are not fundamentally new and can be described in terms of other invariants (for example, it can be shown that $m = 2\tau(K) - \epsilon(K)$).

Fig. 4. A simple figure eight curve contributes two vertical segments.

2.4. $\widehat{\Gamma}(K)$ *and bifiltered complexes*

Readers who are already familiar with knot Floer homology will notice that the object defined above bears little resemblance to the original formulation of the invariant, which takes the form of a $\mathbb{Z}$-graded, $\mathbb{Z}\oplus\mathbb{Z}$-filtered chain complex $CFK^-(K)$ defined up to filtered chain homotopy equivalence. To reassure these readers that the two invariants are in fact equivalent, we pause to briefly describe how the knot Floer complex can be recovered from $\widehat{\Gamma}(K)$. More accurately, we recover the so-called $UV = 0$ quotient of this complex, which records only the horizontal and vertical differentials in $CFK^-(K)$; the stronger invariant $\Gamma(K)$ described in [7], which is $\widehat{\Gamma}(K)$ equipped with some extra decoration, would be required recover the full knot Floer complex.

Let $\widehat{\Gamma} = \widehat{\Gamma}(K)$ be represented by the immersed multicurve $\boldsymbol{\gamma}$ along with decorations as described above, and let μ denote the vertical line through the marked points in $\overline{T}$. We construct a complex $C_{\widehat{\Gamma}}$ over $\mathbb{F}[U, V]/(UV = 0)$ whose generators are the intersection points of $\boldsymbol{\gamma}$ with μ, and whose differential counts immersed bigons for which the left boundary lies on μ and the right boundary is a path in $\boldsymbol{\gamma}$. More precisely, for intersection points x and y in $\boldsymbol{\gamma}\cap\mu$, a bigon from x to y is a homotopy class of maps $f : D^2 \to \overline{T}$ such that $f(-i) = x$, $f(i) = y$, the negative real part of ∂D^2 maps to μ, the positive real part of ∂D^2 maps to $\boldsymbol{\gamma}$, f is an immersion away from $i, -i$, and $f(\partial D^2)$ forms acute corners at x and y. Let $N(x, y)$ denote the mod 2 count of such bigons. We are interested in recording how these bigons cover certain marked points. The marked points of $\overline{T}$ all lie on the line μ; we will push each of these points w slightly off of μ to the right, and add a new marked point z next to each just to the left of μ. For any homotopy class ϕ of maps as above, we define $n_z(\phi)$ and $n_w(\phi)$ to be the multiplicity with which a representative of ϕ covers the z's and w's, respectively. The differential is then given by

$$\partial(x) = \sum_y N(x, y)U^{n_w(x,y)}V^{n_z(x,y)}y.$$

Since we set $UV = 0$ in our coefficient ring, the differential only needs to count bigons that cover either w marked points or z marked points, but not both. We note that to recover the full knot Floer complex, we would need to count bigons covering both types of marked points and we would not set $UV = 0$. However, if we attempt to construct such a complex using only $\widehat{\Gamma}$, ∂^2 may not be zero. To correctly recover the knot Floer complex, we need to take into account some extra decorations in the stronger invariant $\Gamma(K)$ (see [7]).

We set an Alexander grading on the generators of $C_{\widehat{\Gamma}}$, which are intersection points between $\boldsymbol{\gamma}$ and μ, by their height: for $x \in \boldsymbol{\gamma} \cap \mu$, we define $A(x) \in \mathbb{Z}$ to be s if x lies between the marked points at $(0, s - 1/2)$ and $(0, s + 1/2)$. It is clear that if a bigon from x to y covers k marked points of type w, then $A(y) = A(x) + k$, and if it covers k marked points of type z then $A(y) = A(x) - k$; thus if $\partial(x)$ contains a term $U^a V^b y$ where one of a or b vanishes, then $A(y) = A(x) - b + a$.

In addition to the Alexander grading, $C_{\widehat{\Gamma}}$ carries an integer Maslov grading M. This satisfies

$$M(Ux) = M(x) - 2, \quad M(Vx) = M(x), \quad M(\partial x) = M(x) - 1. \tag{2}$$

These relationships determine M as a relative grading on each connected component of $C_{\widehat{\Gamma}}$, since if $U^a V^b y$ appears in ∂x then $M(x) - M(y) = 1 - 2a$. The connected components of $C_{\widehat{\Gamma}}$ correspond directly to the component immersed curves in $\widehat{\Gamma}$. The grading M can be extended to a relative grading on all of $C_{\widehat{\Gamma}}$ by considering bigons between μ and the train track $\boldsymbol{\gamma}_{\text{gr}}$ obtained by including grading arrows with $\boldsymbol{\gamma}$; we require that (2) still holds for these bigons, where running over a grading edge of weight k forward (resp. backward) counts as covering both U and V k times (resp. $-k$ times). That is, if there is a bigon from x to y whose left boundary lies in μ and whose right boundary is a smooth path in $\boldsymbol{\gamma}_{\text{gr}}$ that covers w marked points n_w times and z marked points n_z times, and for which k is the sum of the weights (counted with sign) of all grading edges traversed on the boundary of the bigon traveling from x to y, then

$$M(y) - M(x) = -1 + 2n_w + 2k. \tag{3}$$

We can always assume that all grading arrows in $\widehat{\Gamma}$ lie completely to the right of μ or completely to the left of μ, so to determine the relative Maslov grading it is sufficient to consider bigons that cover only w's or only z's and that include at most one grading arrow. That said, (3) applies for bigons covering both types of marked points, and can in fact be generalized to the following formula for the grading difference between any two generators:

Definition 10. For $x, y \in \boldsymbol{\gamma} \cap \mu$, let P_1 be a path (not necessarily smooth) from x to y in $\boldsymbol{\gamma}_{\text{gr}}$, let P_2 be a path from y to x in μ, and let P be the concatenated path $P_1 P_2$. Then P is a closed path that is smooth apart from right corners at x and y and possibly one or more cusps. Let $\text{rotation}(P)$ denote $\frac{1}{2\pi}$ times the total counterclockwise rotation along the smooth sections of P, let $\text{winding}_w(P)$ denote the net winding number of P around w marked points, and let $\text{weights}(P)$ be the sum of the weights (counted with sign) of all

grading edges traversed by P. Then

$$M(y) - M(x) = -2 \cdot \text{rotation}(P) + 2 \cdot \text{winding}_w(P) + 2 \cdot \text{weights}(P). \quad (4)$$

Note that the mod 2 reduction of the (relative) Maslov grading is determined by the sign of the intersection points in $\boldsymbol{\gamma} \cap \mu$. It is clear that the completeness condition on sets of grading arrows ensures that this relative grading is defined for all generators, and the consistency condition ensures that the relative grading is well defined. Finally, this relative grading can be promoted to an absolute grading by noting that there is a special generator of $C_{\widehat{\Gamma}}$, the first intersection of γ_0 with μ after γ_0 wraps around the cylinder; we set the Maslov grading of this generator to be 0.

As originally defined, the knot Floer complex $CFK^\infty(K)$ is a $\mathbb{Z}$-graded, $\mathbb{Z} \oplus \mathbb{Z}$-filtered chain complex finitely generated over $\mathbb{F}[U, U^{-1}]$. The first filtration is given by negative powers of U, while the second, the Alexander filtration, is recorded separately. The Alexander filtration is determined by an Alexander grading on the generators, together with the fact that multiplication by U lowers the filtration level by 1. It is convenient to add a second formal variable V to keep track of the Alexander filtration, giving rise to a complex generated (with the same generating set) over $\mathbb{F}[U^{\pm 1}, V^{\pm 1}]$, with the two filtrations given by negative powers of U and V, respectively. The original definition is then recovered from this by setting $V = 1$, though the new complex is bigger as many powers of V can be attached to the same element of $CFK^\infty(K)$. To get a complex over $\mathbb{F}[U^{\pm 1}, V^{\pm 1}]$ that is isomorphic to $CFK^\infty(K)$, we would consider only elements $U^a V^b x$ such that $a - b = A(x)$; then a generator x of $CFK^\infty(K)$ would correspond to the element $V^{-A(x)}x$, which is at Alexander filtration level $A(x)$. Terms in the differential that fix the algebraic filtration level (resp. the Alexander filtration level) are referred to as vertical arrows (resp. horizontal arrows); setting $UV = 0$ amounts to counting only horizontal and vertical arrows. $C_{\widehat{\Gamma}}$ as defined above recovers this $UV = 0$ quotient complex.

Example 11. Consider the figure eight knot, whose invariant $\widehat{\Gamma}$ is shown in Figure 3. To recover the knot Floer complex, we would draw the vertical line μ through the marked points, placing basepoints z and w to the left and right of each marked point, and notice that there are five intersections with μ; label these a, b, c, d, e from top to bottom as in Figure 5. The Alexander grading is 1 for a, 0 for b, c, and d, and -1 for e. There are two bigons to the left of μ contributing Vb to ∂a and Ve to ∂c, and there are two bigons to the right of μ contributing Ua to ∂c and Ub to ∂e. The distinguished generator with Maslov grading 0 is d. When we include the grading arrow, there is a bigon on the right of μ from d to a that covers the marked point w once and whose boundary runs over the grading arrow (see Figure 5; note that for clarity we have drawn the pair train track edges that the grading crossover arrow represents, the boundary runs over one of these). The crossover arrow has weight 0, so this bigon implies that $M(a) = 1$. The bigons mentioned previously imply that $M(b) = M(c) = 0$ and $M(e) = -1$.

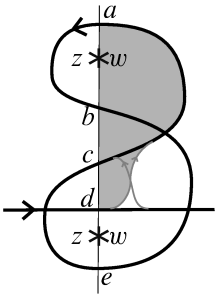

Fig. 5

Remark 12. The discussion above shows that it is fairly straightforward to construct a $UV = 0$ bifiltered complex from a decorated set of immersed curves. The converse, that any $UV = 0$ bifiltered complex can be represented by a decorated set of immersed curves and that this representation is unique in an appropriate sense, is more difficult. This follows from the main Theorem in [8], which proves a related result for type D structures, since the $UV = 0$ quotient of $CFK^\infty(K)$ is equivalent to the bordered Floer invariant of the knot complement $\widehat{CFD}(S^3 \setminus \nu(K))$. See also [7] for a proof that does not pass through bordered Floer homology. $\widehat{\Gamma}(K)$ is defined to be the decorated immersed curve that represents the $UV = 0$ quotient of $CFK^\infty(K)$.

2.5. *Surgery formula*

A key strength of the knot Floer homology package is that there is a simple way to recover the Heegaard Floer homology of any Dehn surgery on a knot K. In particular, $\widehat{HF}(S^3_{p/q}(K))$ can be realized as the intersection Floer homology of the decorated immersed curve $\widehat{\Gamma}(K)$ with lines of slope p/q in the punctured torus $T_\bullet$ or the punctured cylinder $\overline{T}_\bullet$, as we will now further explain (a precise statement is given in Theorem 14).

More precisely, let $\bar{p}(\widehat{\Gamma})$ be the projection of $\widehat{\Gamma}(K)$ to $T_\bullet$; we will see that $\widehat{HF}(S^3_{p/q}(K))$ agrees with intersection Floer homology of $\bar{p}(\widehat{\Gamma})$ with a straight line $\ell_{p,q}$ of slope p/q. By this we mean the homology of a chain complex $CF(\bar{p}(\widehat{\Gamma}), \ell_{p,q})$ generated by intersection points whose differential counts immersed bigons with right boundary on $\bar{p}(\widehat{\Gamma})$ and left boundary on $\ell_{p,q}$. We do not allow bigons to cover the marked point (we indicate this by taking Floer homology in the punctured torus $T_\bullet$ rather than the marked torus T). We count bigons whose boundary includes crossover arrows associated with the local system decoration on $\widehat{\Gamma}$ (see Figure 1 (b)), though it turns out that including these bigons in the differential has no effect on the resulting homology, so in practice the local systems on $\widehat{\Gamma}$ can be ignored. In contrast, we do not count bigons whose boundary runs over a grading arrow so the Maslov decoration has no effect on the differential, but it will be used to define gradings on the resulting complex.

There are two types of grading information on intersection Floer homology. First, $CF(\bar{p}(\widehat{\Gamma}), \ell_{p,q})$ decomposes into spin^c summands, where generators x and y are in the same summand if and only if the loop P formed by concatenating a path from y to x in $\ell_{p,q}$ followed with a (not necessarily smooth) path from x to y in $\bar{p}(\boldsymbol{\gamma}_{\text{gr}})$ (that is, in $\bar{p}(\widehat{\Gamma})$ with grading arrows included) is nullhomologous. This decomposition is easier to understand by lifting to the covering space $\overline{T}_\bullet$, where we take Floer homology of $\widehat{\Gamma}$ with lifts of $\ell_{p,q}$; to recover the same complex we must use multiple different lifts of $\ell_{p,q}$, and the spin^c summands are precisely the Floer homology of $\widehat{\Gamma}$ with any one lift of $\ell_{p,q}$. On each spin^c summand there is also a (relative) Maslov grading, where the grading difference $M(x) - M(y)$ is defined exactly as in Definition 10. The general form of the grading difference can be cumbersome, but in practice it is sufficient to consider bigons that involve at most one grading arrow, possibly with a cusp at one end of the grading arrow.

Lemma 13. *Suppose x and y are connected by an immersed region bounded between $\boldsymbol{\gamma}_{\mathrm{gr}}$ and a lift of $\ell_{p,q}$ that is*

(a) *a bigon not involving a grading arrow,*

(b) *a bigon whose $\boldsymbol{\gamma}_{\mathrm{gr}}$ boundary is a smooth path traversing one grading arrow of weight m, or*

(c) *a cusped bigon whose $\boldsymbol{\gamma}_{\mathrm{gr}}$ boundary traverses one grading arrow of weight m with a single left turning cusp,*

as pictured in Figure 6. Suppose in any case that the region covers k marked points (counted with multiplicity). Then the Maslov grading difference $M(y) - M(x)$ is given by $-1 + 2k$ in case (a)*, $-1 + 2k + 2m$ in case* (b)*, or $2k + 2m$ in case* (c)*.*

Proof. This follows from the general formula for $M(y) - M(x)$; see (4). In cases (a) and (b), the net counterclockwise rotation traversing the closed loop from x to y in $\boldsymbol{\gamma}_{\mathrm{gr}}$ and from y to x in $\ell_{p,q}$ is 2π, but since this includes two right angles at x and y the net rotation along the smooth segments is π; it follows that the term $-2\mathrm{rotation}(P)$ in (4) is -1. In case (c) the extra cusp means that the net rotation along smooth segments in the boundary is 0. In each case the net winding number around the marked points is k, and in cases (b) and (c) the grading arrow contributes m to $\mathrm{weights}(P)$. ∎

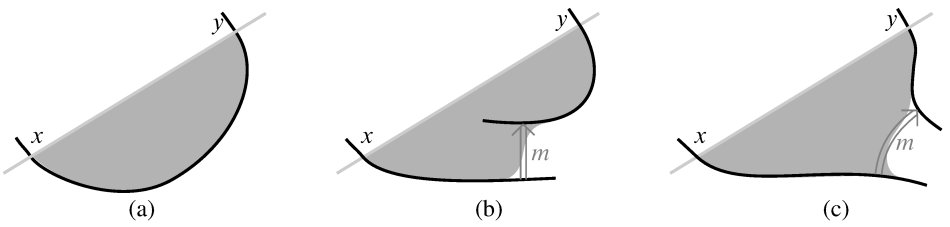

Fig. 6. Three types of regions used to determine the relative Maslov grading on intersection Floer homology. Recall that a crossover arrow consists of a pair of edges in the train track as in Figure 2; the boundaries of the bigons in (b) and (c) use one of these edges.

The following theorem relates $\widehat{HF}$ of surgery on a knot with intersection Floer homology as defined above. Recall that for p/q surgery on a knot $K \subset S^3$ there is a canonical identification of $\mathrm{Spin}^c(S^3_{p/q}(K))$ with $\mathbb{Z}/p\mathbb{Z}$.

Theorem 14 (Surgery formula). *Consider a knot $K \in S^3$ and p, q relatively prime with $p > 0$. Fix a small $\epsilon > 0$ (in particular, $\epsilon < 1/q$), and for each $i \in \mathbb{Z}/p\mathbb{Z}$ let $\ell^i_{p,q}$ be a straight line in the punctured cylinder $\overline{T}_\bullet$ of slope p/q that passes through the point $(0, -\frac{1}{2} + \frac{i}{q} + \epsilon)$. Then*

$$\widehat{HF}(S^3_{p/q}, i) \cong HF(\widehat{\Gamma}(K), \ell^i_{p,q})$$

as relatively graded vector spaces, where the right hand side refers to intersection Floer homology in $\overline{T}_\bullet$.

Proof. This is a special case of a much more general gluing formula for bordered Floer invariants ([8, Theorem 2] without gradings and [9, Theorem 5] with gradings). Indeed, $\widehat{\Gamma}(K)$ is precisely the invariant $\widehat{HF}(M)$ associated with the knot complement $M = S^3 \setminus \nu(K)$, the invariant $\widehat{HF}(D^2 \times S^1)$ is simply the meridian $\partial D^2 \times S^1$, and p/q-Dehn surgery corresponds to gluing $D^2 \times S^1$ to M by a map taking the meridian to a line of slope p/q. ∎

A direct consequence is that $\mathrm{rk}(\widehat{HF}(S^3_{p/q}(K)))$ is given by the minimal intersection in $T_\bullet$ of $\bar{p}(\widehat{\Gamma})$ and $\ell_{p,q}$. This is because all bigons not covering a puncture can be removed by pulling $\bar{p}(\widehat{\Gamma})$ tight (here we need that no component of $\bar{p}(\widehat{\Gamma})$ is parallel to $\ell_{p,q}$ to ensure admissibility, but this is clear if $p \neq 0$). Similarly, $\mathrm{rk}(\widehat{HF}(S^3_{p/q}(K), i))$ is the minimal intersection in $\overline{T}_\bullet$ of $\widehat{\Gamma}$ and $\ell^i_{p,q}$. We remark that when $\widehat{\Gamma}$ is pulled tight as described in Section 2.3, it automatically has minimal intersection with each $\ell^i_{p,q}$.

Recall that if $\widehat{\Gamma}$ is pulled tight, then outside of a neighborhood of the punctures it consists of a collection of some number n of length 1 vertical segments and a single non-vertical segment of slope m. Then we have the following expression for $\mathrm{rk}(\widehat{HF}(S^3_{p/q}(K)))$ (compare [19, Proposition 9.5]):

Proposition 15. *With the integers m and n defined as above,*

$$\mathrm{rk}(\widehat{HF}(S^3_{p/q}(K))) = |p - mq| + n|q|.$$

Proof. A line of slope p/q in T intersects the vertical line through the marked point $|q|$ times, so there are $n|q|$ intersection points coming from vertical segments in $\bar{p}(\widehat{\Gamma})$. The remaining intersection points come from intersections with the segment of slope m, and the number of such intersections is the distance between the slopes p/q and $m/1$, namely $\left|\det \left(\begin{smallmatrix} p & q \\ m & 1 \end{smallmatrix}\right)\right|$. ∎

Remark 16. The key idea in the proof of [8, Theorem 2] (in the special case required for Theorem 14) is to perturb both $\widehat{\Gamma}(K)$ and $\ell^i_{p,q}$ into a special form so that the intersection Floer chain complex can be directly identified with (one spin^c component of) the box tensor product of two bordered Floer invariants, the homology of which is known to agree with $\widehat{HF}(S^3_{p/q}, i)$. It is also possible to prove Theorem 14 without passing through bordered Floer homology. This is accomplished by perturbing $\ell^i_{p,q}$ so that the intersection Floer chain complex $CF(\widehat{\Gamma}(K), \ell^i_{p,q})$ is identified with the complex $\widehat{\mathbb{X}}_i$ in the mapping cone formula [19], whose homology is also known to give $\widehat{HF}(S^3_{p/q}, i)$. This identification was shown for large integer surgery in [9, Proposition 70], and the full proof will appear in [7].

While the two proofs are similar in spirit, this second proof has a few advantages since the mapping cone formula carries some information not available with bordered Floer homology. For example, the mapping cone formula recovers the absolute grading on $\widehat{HF}(S^3_{p/q}, i)$ while bordered Floer homology can only give the relative grading. In addition, the identification of the Floer chain complex with the mapping cone formula

can be generalized to one involving $\mathbb{X}_i^+$ instead of $\widehat{\mathbb{X}}_i$, so a version of Theorem 14 holds for $+$ type invariants [7]. We will not need the absolute grading or $+$ type invariants in the present paper, but we will make use of the identification mentioned in Remark 16 in one small way in Proposition 17 below; namely, we will use the fact that the subset of generators of $HF(\widehat{\Gamma}(K), \ell^i_{p,q})$ arising from any one curve component γ_i in $\widehat{\Gamma}(K)$ can be identified with the subset of generators of $S^3_{p,q}(K)$ arising from the mapping cone formula applied to direct summand of $CFK^\infty(K)$ corresponding to γ_i.

We will be interested in a special class of knots for which the bifiltered complex $CFK^\infty(K)$ has a direct summand that looks like $CFK^\infty(U)$, where U is the unknot. That is, we require that for some choice of basis $CFK^\infty(K)$ has a generator with no differentials in or out. In this case we will say that $CFK^\infty(K)$ *has an isolated generator*. Note that $CFK^\infty(K)$ having an isolated generator implies that the curve γ_0 in $\widehat{\Gamma}(K)$ is homotopic to the horizontal curve wrapping around the cylinder once, but the latter condition is slightly weaker since γ_0 does not see diagonal arrows in $CFK^\infty(K)$ (giving an immersed curve condition equivalent to having an isolated generator would require the stronger invariant $\Gamma(K)$ and amounts to further imposing that γ_0 is not connected to any other γ_i by the additional decorations in $\Gamma(K)$).

Recall that for a 3-manifold Y with spin^c structure $\mathfrak{s}$, the d-invariant, or correction term, $d(Y,\mathfrak{s})$ is defined as the minimum absolute grading of an element of the image of $HF^\infty(Y,\mathfrak{s})$ in $HF^+(Y,\mathfrak{s})$. Understanding d-invariants usually requires working the $+$ flavor of invariants; there is always a generator in $\widehat{HF}(Y,\mathfrak{s})$ whose absolute grading is the d-invariant, but without knowing the U-module structure on HF^+ we generally have no way of knowing which generator gives the d-invariant (unless, of course, there is only one generator in $\widehat{HF}(Y,\mathfrak{s})$). However, if Y is (nonzero) surgery on a knot K for which $CFK^\infty(K)$ has an isolated generator, there is an obvious choice for a distinguished generator in each spin^c structure and indeed this generator gives the d-invariant.

Proposition 17. *Suppose K is a knot for which $CFK^\infty(K)$ has an isolated generator, and in particular the distinguished curve γ_0 of $\widehat{\Gamma}(K)$ is horizontal. Then for each $i \in \mathbb{Z}/p\mathbb{Z}$, the absolute grading of the generator of $\widehat{HF}(S^3_{p/q}(K), i) \cong HF(\widehat{\Gamma}(K), \ell^i_{p,q})$ corresponding to the unique intersection point of γ_0 with $\ell^i_{p,q}$ is $d(S^3_{p/q}(K), i)$.*

Proof. In the mapping cone formula, the direct summands of $CFK^\infty(K)$ give rise to direct summands for the mapping cone $\mathbb{X}_i$, and it is clear that to compute the d-invariant it is sufficient to consider only the unique nonacyclic summand of $CFK^\infty(K)$ and the corresponding summand of the mapping cone. When $CFK^\infty(K)$ has an isolated generator, the homology of this summand has rank 1, so the d-invariant must be the grading of its only generator. We now appeal not just to Theorem 14 but also to the identification of $CF(\widehat{\Gamma}(K), \ell^i_{p,q})$ with $\widehat{\mathbb{X}}_i$ mentioned in Remark 16. The direct summands of $CFK^\infty(K)$ correspond to the curve components of $\widehat{\Gamma}(K)$, with the nonacyclic summand corresponding to γ_0, and so the relevant summand of $\widehat{\mathbb{X}}_i$ is identified with the intersection Floer complex of γ_0 with $\ell^i_{p,q}$. Thus the grading of the unique generator of $HF(\gamma_0, \ell^i_{p,q})$ is $d(S^3_{p/q}(K), i)$. ∎

This result is not at all surprising, but it does require the mapping cone formula proof of Theorem 14 since the bordered Floer approach gives no way of confirming that the obvious distinguished summand $HF(\gamma_0, \ell^i_{p,q})$ should capture the d-invariant. We remark that this use of Remark 16 in the proof of Proposition 17 is the only essential dependence of the present paper on [7].

3. Obstructing truly cosmetic surgeries

We now turn to a brief survey of some past results on which the arguments in the next section build. The first observation is that, since $H_1(S^3_{p,q}(K)) \cong \mathbb{Z}/p\mathbb{Z}$, any pair of cosmetic surgery slopes must have the same numerator. The next constraint is a condition on the Alexander polynomial of K proved by Boyer and Lines:

Theorem 18 ([1, Proposition 5.1]). *If K admits a truly cosmetic surgery, then $\Delta''_K(1) = 0$.*

This result is a consequence of surgery formulas for the Casson–Walker invariant λ and the Casson–Gordon invariant $\boldsymbol{\tau}$:

$$\lambda(S^3_{p/q}(K)) = \lambda(L(p,q)) + \frac{q}{2p}\Delta''_K(1), \tag{5}$$

$$\boldsymbol{\tau}(S^3_{p/q}(K)) = \boldsymbol{\tau}(L(p,q)) - \sigma(K,p), \tag{6}$$

where $\sigma(K,p) = \sum_{r=0}^{p-1} \sigma_K(e^{2i\pi r/p})$ does not depend on q. If p/q and p/q' are truly cosmetic surgery slopes, (6) implies that $\boldsymbol{\tau}(L(p,q)) = \boldsymbol{\tau}(L(p,q'))$. For a lens space, $\boldsymbol{\tau}(L(p,q))$ is a constant multiple of $p\lambda(L(p,q))$, so in fact $\lambda(L(p,q)) = \lambda(L(p,q'))$. Then (5) implies that either $q = q'$ or $\Delta''_K(1) = 0$.

Heegaard Floer homology entered the story when Ozsváth and Szabó constructed a surgery formula in terms of knot Floer homology [19] and used it to prove the following proposition. As a demonstration of the machinery that will be used in this paper, we present a proof that is essentially equivalent to the one in [19] but is reframed in the language of the immersed curve surgery formula.

Proposition 19 ([19, Theorem 1.5]). *Suppose $S^3_{p/q_1}(K) \cong \pm S^3_{p/q_2}(K)$ with $q_1 \neq q_2$. Either q_1 and q_2 have opposite signs or $S^3_{p/q_1}(K)$ is an L-space.*

Proof. We must have $\operatorname{rk}(\widehat{HF}(S^3_{p/q_1}(K))) = \operatorname{rk}(\widehat{HF}(S^3_{p/q_2}(K)))$. By Proposition 15,

$$|p - mq_1| + n|q_1| = |p - mq_2| + n|q_2|, \tag{7}$$

where m is the slope of the nonvertical segment in $\widehat{\Gamma}(K)$ and n is the number of vertical segments. By taking the mirror of K if necessary, we may assume without loss of generality that $m \geq 0$. First suppose that q_1 and q_2 are both negative or that they are both positive and greater than p/m. In either case, (7) simplifies to $(m+n)q_1 = (m+n)q_2$. Since $m + n > 0$ for a nontrivial knot, this implies that $q_1 = q_2$. Next suppose that q_1 and q_2 are both positive and smaller than p/k; in this case (7) simplifies to $(n-m)q_1 =$

$(n-m)q_2$. If $q_1 \neq q_2$, we must have $n = m$, which implies that K is an L-space knot, and since $p/q_i > m = 2g(K) - 1$, the result of either surgery is an L-space. Finally, suppose that q_1 and q_2 are both positive, with $mq_1 < p$ and $mq_2 > p$; then (7) becomes $p - mq_1 + nq_1 = mq_2 - p + nq_2$. This implies

$$n(q_2 - q_1) = 2p - m(q_1 + q_2) < 2mq_2 - m(q_1 + q_2) = m(q_2 - q_1).$$

This is a contradiction, since $n \geq m$. ∎

In the case of truly cosmetic surgery on a knot K with Seifert genus 1, Ozsváth and Szabó in fact showed that the surgery must be an L-space [19, Theorem 1.4]. Wang ruled out this possibility, implying that the cosmetic surgery conjecture holds for all genus 1 knots [23]; we will give a new proof of this fact in Section 4 (see Corollary 26). Wu later ruled out the possibility that truly cosmetic surgeries are L-spaces for arbitrary knots [24] by observing that the restrictions on the Alexander polynomial of an L-space knot given in [18] imply that $\Delta_K''(1) \neq 0$ and applying Theorem 18. Thus truly cosmetic surgery slopes have opposite sign.

A significant advancement came in the following result of Ni and Wu:

Theorem 20 ([15, Theorem 1.2]). *Suppose $S^3_{p/q} \cong S^3_{p/q'}$ with $q' \neq q$. Then*

(i) $\tau(K) = 0$, *where τ is the Ozsváth–Szabó concordance invariant;*

(ii) $q' = -q$; *and*

(iii) $q^2 \equiv -1 \pmod{p}$.

The key ingredient here was a surgery formula for the d-invariants in Heegaard Floer homology [15, Proposition 1.6]. A consequence of the surgery formula is that for $p/q > 0$, the d-invariants of $S^3_{p/q}(K)$ are less than or equal to the corresponding d-invariants of $S^3_{p/q}(U) = L(p,q)$, with equality holding for all spin^c structures if and only if $V_0(K) = H_0(K) = 0$, where V_0 and H_0 are integer invariants related to certain maps in the rational surgery formula. For $p/q' < 0$, the same relationship holds with the inequality reversed. Let $d(Y)$ denote $\sum_{\mathfrak{s} \in \mathrm{Spin}^c(Y)} d(Y, \mathfrak{s})$. For a lens space, $d(L(p,q))$ is a constant multiple of the Casson–Walker invariant $\lambda(L(p,q))$, and it was already noted that for a truly cosmetic surgery equations (5) and (6) imply that $\lambda(L(p,q)) = \lambda(L(p,q'))$. Thus

$$d(S^3_{p/q}) \leq d(L(p,q)) = d(L(p,q')) \leq d(S^3_{p/q'}).$$

For a truly cosmetic surgery equality must hold, so $V_0(K) = H_0(K) = 0$. This in particular implies (i), and then by Proposition 15, $\mathrm{rk}(\widehat{HF}(S^3_{p/q}(K)))$ is a linear function of $|q|$, which implies (ii). Finally, (iii) follows from the fact that $d(L(p,q)) = d(L(p,-q)) = -d(L(p,q))$, and an explicit formula for $\lambda(L(p,q))$ showing that $\lambda(L(p,q)) = 0$ if and only if $q^2 \equiv -1 \pmod{p}$.

In fact, the first conclusion is slightly understated, since the proof really shows that $V_0(K) = H_0(K) = 0$ [15, Theorem 2.5], and this is strictly stronger than $\tau(K)$ being zero. Hom showed that when this condition holds then $CFK^\infty(K)$ has an isolated generator [12, Proposition 3.11]. Recall that by this we mean that for some choice of basis,

$CFK^\infty(K)$ has a single generator with no differentials in or out. Hom's paper also shows that the bifiltered chain complex $CFK^\infty(K)$, taken up to filtered chain homotopy equivalence and up to adding and removing acyclic summands, is a concordance invariant from which all known Heegaard Floer concordance invariants can be derived; having an isolated generator is equivalent to this concordance invariant being trivial. In the language of immersed curves, $CFK^\infty(K)$ having an isolated generator implies that γ_0 is the horizontal curve wrapping around the cylinder once; this in turn is equivalent to $\epsilon(K)$ being 0 and implies $\tau(K) = 0$. To summarize, we have the following implications:

$$CFK^\infty(K) \text{ has isolated genarator} \underset{\not\Longleftarrow}{\Longrightarrow} \gamma_0 \text{ is horizontal} \iff \epsilon(K) = 0 \underset{\not\Longleftarrow}{\Longrightarrow} \tau(K) = 0.$$

Thus, Ni and Wu really proved the following:

Theorem 21 ([15, Theorem 1.2, enhanced]). *Suppose $S^3_{p/q} \cong S^3_{p/q'}$ with $q' \neq q$. Then*

(i) *$CFK^\infty(K)$ has an isolated generator, in particular, $\epsilon(K) = \tau(K) = 0$;*

(ii) *$q' = -q$; and*

(iii) *$q^2 \equiv -1 \pmod{p}$.*

It makes sense that the original theorem was stated in terms of τ only, as ϵ had not been defined at that time and the condition that $V_0 = H_0 = 0$ or that CFK^∞ has an isolated generator makes for a more cumbersome statement. However, this means that some implications of Ni and Wu's work, which has already found many wonderful applications, have been overlooked. For example, the following result follows immediately from Theorem 21 and a cabling formula of Hom [11, Theorem 2], which says that ϵ of a cable is never zero:

Corollary 22. *The cosmetic surgery conjecture holds for any nontrivial cable of a knot in S^3.*

This result was recently proved in [22] using Theorem 20 and Hom's cabling formula for τ [11, Theorem 1] to rule out many cases, but other methods were needed to deal with cables for which $\tau = 0$.

We end this section with a technical result that will be required later, related to one used by Ni and Wu in the proof of Theorem 20. Recall that part (iii) of that theorem follows from the fact that if p/q is a truly cosmetic surgery slope then the sum of all the d-invariants of $L(p,q)$ must be 0. More precisely, there is an explicit formula for this sum of d-invariants [20, Lemmas 2.2 and 4.3]:

$$d(L(p,q)) := \sum_{i=0}^{p-1} d(L(p,q),i) = p\lambda(L(p,q)) = -\frac{1}{12}\left[q + q' + p\sum_{i=1}^{n}(a_i - 3)\right], \tag{8}$$

where q' is the unique integer $0 < q' < p$ with $qq' \equiv 1 \pmod{p}$ and $[a_1, \ldots, a_n]$ is the Hirzebruch–Jung continued fraction expansion for p/q. If this sum is 0, then considering

the term in the brackets modulo p implies that $q \equiv -q' \pmod{p}$ and Theorem 20 (iii) follows. We will at times be interested in only the first q d-invariants of $L(p,q)$; below we show that when the sum of all p d-invariants of $L(p,q)$ is 0, then the sum of the first q of them is nonzero.

Lemma 23. *For $p > q > 0$ relatively prime and $q^2 \equiv -1 \pmod{p}$,*

$$\sum_{i=0}^{q-1} d(L(p,q),i) \neq 0.$$

Proof. We will show that the sum is nonzero modulo $1/12$. We use the recursive formula for d-invariants of $L(p,q)$ given by Ozsváth and Szabó [16, Proposition 4.8]:

$$d(L(p,q),i) = -\frac{1}{4} + \frac{(2i+1-p-q)^2}{4pq} - d(L(q,r),j),$$

where r and j are the mod q reductions of p and i, respectively. In particular,

$$\sum_{i=0}^{q-1} d(L(p,q),i) = -\sum_{i=0}^{q-1} \frac{1}{4} + \sum_{i=0}^{q-1} \frac{(2i+1-p-q)^2}{4pq} - \sum_{i=0}^{q-1} d(L(q,r),i).$$

The third sum on the right hand side is simply $d(L(q,r))$, and it it is easy to see from (8) that this is an integer multiple of $1/12$. The first sum on the right hand side, which evaluates to $q/4$, is also a multiple of $1/12$, so it is enough to check that the second sum is not. We have

$$\begin{aligned}\sum_{i=0}^{q-1}(2i+1-p-q)^2 &= \sum_{i=0}^{q-1}\big(p^2 + 2p(q-1-2i) + (q-1-2i)^2\big)\\ &= qp^2 + 2p\sum_{i=0}^{q-1}(q-1-2i) + \sum_{i=0}^{q-1}(q-1-2i)^2.\end{aligned}$$

The second term in the expression on the right is 0, since the summands run evenly from $q-1$ to $-(q-1)$. The final term is twice the sum of the first $q/2$ odd squares if q is even, or twice the sum of the first $(q-1)/2$ even squares if q is odd; in either case, the sum evaluates to $q(q-1)(q+1)/3$. Thus we need to show that

$$\frac{qp^2 + q(q^2-1)/3}{4pq} = \frac{p}{4} + \frac{q^2-1}{12p}$$

is not a multiple of $1/12$. The first term clearly is, but the second term is not as long as $q^2 \not\equiv 1 \pmod{p}$. This holds in particular when $q^2 \equiv -1 \pmod{p}$, unless $p = 2$. We complete the proof by directly checking the case $p = 2$, $q = 1$: the claim holds since $d(L(2,1),0) = 1/4 \neq 0$. ∎

4. New obstructions

Throughout this section we fix a knot K, and let $\widehat{\Gamma} = \widehat{\Gamma}(K)$ with underlying set $\boldsymbol{\gamma}$ of immersed curves. We will assume that $CFK^\infty(K)$ has an isolated vertex, which by Theorem 21 is necessary for K to admit a truly cosmetic surgery. In particular, this means that the distinguished curve γ_0 in $\boldsymbol{\gamma}$ is horizontal. Theorem 21 also says that any pair of cosmetic surgery slopes are opposite, so we fix $p, q > 0$ relatively prime, and set $Y_+ = S^3_{p/q}(K)$ and $Y_- = S^3_{-p/q}(K)$. Our goal is to obstruct Y_+ and Y_- from being orientation preserving diffeomorphic by finding conditions under which $\widehat{HF}(Y_+)$ and $\widehat{HF}(Y_-)$ are not isomorphic as graded vector spaces.

The results in the previous section primarily make use of the Casson–Walker and Casson–Gordon invariants, the total rank of $\widehat{HF}$, and the d-invariants, which can be viewed as the Maslov grading of one special generator of $\widehat{HF}$ for each spin^c structure. To extract more information and produce new obstructions, we will need to use the Maslov grading of all generators. In particular, the set of gradings of all generators of $\widehat{HF}$ is an invariant, as is the partitioning of this set into subsets according to spin^c structures. To avoid working with absolute gradings, we define

$$M_{\mathrm{rel}}(x) = M(x) - d(Y, \mathfrak{s})$$

for x in $\widehat{HF}(Y, \mathfrak{s})$. Since γ_0 is horizontal, for each $i \in \mathbb{Z}/p\mathbb{Z} \cong \mathrm{Spin}^c(Y_\pm)$ there is a distinguished generator in $\widehat{HF}(Y_\pm, i)$ coming from the unique intersection point of γ_0 with $\ell^i_{p,\pm q}$; we will denote this generator x^i_0. By Proposition 17 the absolute grading of x^i_0 is $d(Y_\pm, i)$. Thus for Y_+ and Y_-, $M_{\mathrm{rel}}(x)$ is simply the Maslov grading relative to the distinguished generator in the same spin^c structure, i.e. $M_{\mathrm{rel}}(x) = M(x) - M(x^i_0)$. We will consider the following multisets (that is, sets with repetition allowed) of relative gradings:

$$M_{\mathrm{rel}}(Y) = \{M_{\mathrm{rel}}(x) \mid x \text{ a generator of } \widehat{HF}(Y)\},$$
$$M_{\mathrm{rel}}(Y, \mathfrak{s}) = \{M_{\mathrm{rel}}(x) \mid x \text{ a generator of } \widehat{HF}(Y, \mathfrak{s})\}.$$

These are invariants of Y and the pair $(Y, \mathfrak{s})$, respectively. In particular, if $Y_+ \cong Y_-$ then the sets $M_{\mathrm{rel}}(Y_+)$ and $M_{\mathrm{rel}}(Y_-)$ agree. Moreover, there is some permutation σ on $\mathbb{Z}/p\mathbb{Z}$ such that $M_{\mathrm{rel}}(Y_+, i) = M_{\mathrm{rel}}(Y_-, \sigma(i))$. We will at times refer to the sum of all elements in these sets, which we denote $\Sigma M_{\mathrm{rel}}(Y)$ and $\Sigma M_{\mathrm{rel}}(Y, \mathfrak{s})$, respectively.

Remark 24. Both Spin$^c(Y_+)$ and Spin$^c(Y_-)$ can be identified with $\mathbb{Z}/p\mathbb{Z}$ in a way that is canonical *given the surgery description*, but this identification is not an invariant of the manifold. Thus even if $Y_+ \cong Y_-$, the ith spin^c structure of Y_+ need not agree with the ith spin^c structure of Y_-; this is why the permutation σ is required above.

It is easy to see that the ranks of $\widehat{HF}(Y_+)$ and $\widehat{HF}(Y_-)$ agree. Indeed, since γ_0 is horizontal, the slope m of the nonvertical segment in $\widehat{\Gamma}$ is 0, so by Proposition 15,

$$\mathrm{rk}(\widehat{HF}(Y_+)) = p + n|q| = p + n|-q| = \mathrm{rk}(\widehat{HF}(Y_-)).$$

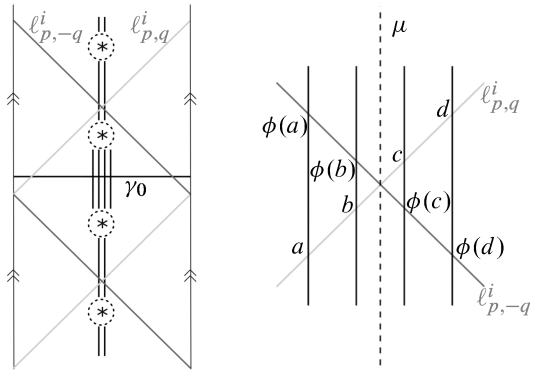

Fig. 7. The intersection of $\widehat{\Gamma}$ (black) with a line $\ell^i_{p,q}$ (light gray) and its vertical reflection $\ell^i_{p,-q}$ (dark gray). We assume that γ_0 is a simple horizontal curve, and that outside of a neighborhood of the marked points the remaining immersed curves in $\widehat{\Gamma}$ consist of parallel vertical segments. The map ϕ_i takes an intersection of $\ell_{p,q}$ with a vertical segment to the nearby intersection of $\ell^i_{p,-q}$ with that segment, as pictured to the right.

Our main strategy for studying the sets of gradings described above is to define for each $i \in \mathbb{Z}/p\mathbb{Z}$ a particular map $\phi_i\colon \widehat{HF}(Y_+, i) \to \widehat{HF}(Y_-, i)$ which is an isomorphism of ungraded vector spaces. In other words, ϕ_i gives a one-to-one correspondence between generators of $\widehat{HF}(Y_+, i)$ and generators of $\widehat{HF}(Y_-, i)$. This correspondence does not preserve the relative grading M_{rel}, even when $\widehat{HF}(Y_+)$ and $\widehat{HF}(Y_-)$ are isomorphic as graded vector spaces; however, we will be able to say explicitly how M_{rel} changes under ϕ_i and we can use this to determine if the sets of gradings defined above are fixed. Combining these maps for all i gives an (ungraded) isomorphism $\phi\colon \widehat{HF}(Y_+) \to \widehat{HF}(Y_-)$.

We will assume that $\widehat{\Gamma}$ has the form described in Section 2.3 and shown in Figure 4. That is, we assume the curve is pulled tight, noting that outside of a neighborhood of the marked points each curve γ_i with $i \neq 0$ consists of a collection of roughly vertical segments, and we perturb the curves slightly so that these are in fact parallel vertical segments (see Figure 7). The endpoints of these vertical arcs are connected in some way within the neighborhoods of the marked points, but this information will not be relevant to us. Recall that n_s denotes the number of these vertical segments at height s, and $n = \sum_{s=-\infty}^{\infty} n_s$ is the total number of vertical segments. When $\widehat{\Gamma}$ is pulled tight in this way, it is clear that it intersects minimally with $\ell^i_{p,q}$ and $\ell^i_{p,-q}$, so we may view $\widehat{HF}(Y_\pm, i) \cong HF(\widehat{\Gamma}, \ell^i_{p,\pm q})$ as generated by $\widehat{\Gamma} \cap \ell^i_{p,\pm q}$.

We will now define ϕ_i by describing a one-to-one correspondence between $\widehat{\Gamma} \cap \ell^i_{p,q}$ and $\widehat{\Gamma} \cap \ell^i_{p,-q}$. Note that $\ell^i_{p,-q}$ is the reflection across μ of $\ell^i_{p,q}$. In particular, each intersection of $\ell^i_{p,q}$ with μ is also an intersection of $\ell^i_{p,-q}$ with μ. The vertical segments are arbitrarily close to μ, so each intersection of $\ell^i_{p,q}$ with a vertical segment can be uniquely identified with an intersection point in $\mu \cap \ell^i_{p,q}$, namely the nearest such intersection point. Conversely, for each point in $\mu \cap \ell^i_{p,q}$ and for each vertical segment at the same height as that point, there is exactly one nearby intersection of the vertical segment

with $\ell^i_{p,q}$. The same is true for intersections between a vertical segment and $\ell^i_{p,-q}$. If $x \in \widehat{\Gamma} \cap \ell^i_{p,q}$ lies on a vertical segment in $\widehat{\Gamma}$, then we define $\phi_i(x)$ to be the point on $x \in \widehat{\Gamma} \cap \ell^i_{p,-q}$ that lies on the same vertical segment in $\widehat{\Gamma}$ and corresponds to the same point in $\mu \cap \ell^i_{p,q} = \mu \cap \ell^i_{p,-q}$; see the right side of Figure 7 depicting a neighborhood of one point in $\mu \cap \ell^i_{p,\pm q}$. If $x \in \widehat{\Gamma} \cap \ell^i_{p,q}$ does not lie on a vertical segment in $\widehat{\Gamma}$ then it is the unique point in $\gamma_0 \cap \ell^i_{p,q}$, corresponding to the distinguished generator of $\widehat{HF}(Y_+, i)$; in this case we define $\phi_i(x)$ to be the unique point in $\gamma_0 \cap \ell^i_{p,-q}$, so ϕ_i identifies the distinguished generators.

We often need to distinguish the generators in $\widehat{HF}(Y_\pm)$ that come from intersections on vertical segments of $\widehat{\Gamma}$ from the distinguished generators coming from intersections on γ_0, which we denote by x^i_0 in the ith spin^c structure. Mimicking the standard notation for HF^+, we will use $\widehat{HF}_{\text{red}}(Y_\pm, i)$ to denote the summand of $\widehat{HF}(Y_\pm, i)$ obtained by removing the generator x^i_0, and $\widehat{HF}_{\text{red}}(Y_\pm)$ to denote $\bigoplus_{i \in \mathbb{Z}/p\mathbb{Z}} \widehat{HF}_{\text{red}}(Y_\pm, i)$. Of course, reduced Floer homology in the hat setting does not make sense in general, but in this case, where we know that the distinguished generator x^i_0 gives the d-invariant and corresponds to the tower in HF^+, the analogy is appropriate.

For each generator x of $\widehat{HF}_{\text{red}}(Y_+)$, we are interested in computing both $M_{\text{rel}}(x)$ and $M_{\text{rel}}(\phi(x))$ (for the distinguished generators x^i_0, both quantities are 0 by definition). The generator x corresponds to an intersection point between $\ell^i_{p,q}$ for some i and a vertical segment of $\widehat{\Gamma}$. There are two integers we will associate with such an intersection point x. First, let $A(x)$ denote the height of the relevant vertical segment. Second, after a slight perturbation we can assume that the vertical segment containing x lies exactly on the vertical line μ through the marked points of $\overline{T}$ and that x can be viewed as an intersection point of μ and $\ell^i_{p,q}$; let $k(x)$ denote the number of marked points, counted with multiplicity, in the interior of the triangle formed by μ, $\ell^i_{p,q}$, and γ_0. It is easiest to picture this triangle in the covering space $\widetilde{T}$, as shown in Figure 8. Using these quantities, we can compute the effect that the map ϕ has on the relative grading of x.

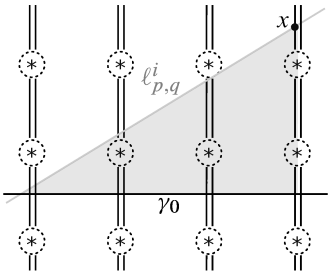

Fig. 8. The generator x corresponding to the marked intersection point has $A(x) = 2$ since the intersection lies on a vertical segment of height 2, and $k(x) = 2$ since the interior of the shaded triangle covers two marked points.

Proposition 25. *For a generator $x \in \widehat{HF}(Y_+, i)$ corresponding to an intersection of $\ell^i_{p/q}$ with a vertical segment in $\widehat{\Gamma}$, let $A(x)$ and $k(x)$ be the quantities defined above. Then*

$$\Delta_{\mathrm{rel}}(x) := M_{\mathrm{rel}}(\phi(x)) - M_{\mathrm{rel}}(x) = 1 - 2|A(x)| - 4k(x).$$

Proof. We will assume that $A(x) \geq 0$; if $A(x) < 0$, the proof is exactly the same with all pictures rotated 180 degrees and $A(x)$ replaced with $|A(x)|$. We will only work with γ_0 and the vertical segment containing x, and will ignore the rest of $\widehat{\Gamma}$. Up to perturbing $\widehat{\Gamma}$ we may assume that the vertical segment in question lies exactly on μ, so that the points $x \in \widehat{\Gamma} \cap \ell^i_{p,q}$ and $\phi(x) \in \widehat{\Gamma} \cap \ell^i_{p,-q}$ coincide; this means that $\ell^i_{p,q}$, $\ell^i_{p,-q}$, and γ_0 form a triangle. In order to compute M_{rel}, we need the grading decoration on $\widehat{\Gamma}$. We can assume that the set of grading arrows contains an arrow that lies on the right side of μ and connects γ_0 to the bottom end of the vertical segment; this grading arrow carries some integer weight m. There are two cases to consider, depending on whether the vertical segment containing x is oriented up or down (note that we always assume γ_0 is oriented rightward). The first case is pictured on the left of Figure 9; in this case the grading arrow goes from γ_0 to the right side of the vertical segment. There is a bigon from x^i_0 to x, shaded dark gray in the figure, which has no cusps, covers $k(x) + A(x)$ punctures, and whose boundary runs over the grading arrow labeled by m. By Lemma 13 (b),

$$M_{\mathrm{rel}}(x) = M(x) - M(x^i_0) = -1 + 2k(x) + 2A(x) + 2m.$$

The complement of this region within the triangle formed by $\ell^i_{p,q}$, $\ell^i_{p,-q}$, and γ_0, shaded light gray in Figure 9, is a cusped bigon from $\phi(x)$ to $\phi(x^i_0)$. This bigon covers $k(x)$, and its boundary runs over the grading arrow backwards and has a single cusp, at the tail of the grading arrow. It follows from Lemma 13 (c) that

$$M_{\mathrm{rel}}(\phi(x)) = M(\phi(x)) - M(\phi(x^i_0) = -2k(x) + 2m.$$

Thus $\Delta_{\mathrm{rel}}(x) = M_{\mathrm{rel}}(\phi(x)) - M_{\mathrm{rel}}(x) = 1 - 2A(x) - 4k(x)$, as desired. Note that the label of the grading arrow cancels out and does not end up affecting $\Delta_{\mathrm{rel}}(x)$.

If the vertical segment is oriented downward, the grading arrow must go to the left side of the vertical segment to be consistent with the orientations. The right side of Figure 9 shows the modified grading arrow we will use. The only difference is that the boundary of the dark gray bigon from x^i_0 to x now has one cusp while the bigon from $\phi(x)$ to $\phi(x^i_0)$ can be drawn with no cusps. This change adds 1 to $M_{\mathrm{rel}}(x)$ and also adds 1 to $M_{\mathrm{rel}}(\phi(x))$, so it does not affect $\Delta_{\mathrm{rel}}(x)$. ∎

Note that the triangle formed by $\ell^i_{p,q}$, $\ell^i_{p,-q}$, and γ_0 covers $|A(x)| + 2k(x)$ punctures, so proposition says that $\Delta_{\mathrm{rel}}(x)$ is 1 minus twice the number of punctures covered by this triangle. This number of punctures is nonnegative. Moreover, $k(x) = 0$ if $A(x) = 0$; it follows that $\Delta_{\mathrm{rel}}(x) = 1$ if and only if $A(x) = 0$, and $\Delta_{\mathrm{rel}}(x) < 0$ otherwise. An immediate corollary of this is a reproof of a result of Wang:

Corollary 26 ([23, Theorem 1.3]). *If $g(K) = 1$, then K does not admit truly cosmetic surgeries.*

Proof. If K admits a truly cosmetic surgery, then by Theorem 21 we may assume that the slopes are opposite and $CFK^\infty(K)$ has an isolated generator. Let $Y_+ = S^3_{p/q}(K)$ and $Y_- = S^3_{-p/q}(K)$, and define $\phi : \widehat{HF}(Y_+) \to \widehat{HF}(Y_-)$ as above. If $g(K) = 1$, then all vertical segments in $\widehat{\Gamma}$ are at height 0. It follows that every intersection x of any line of slope p/q with a vertical segment of $\widehat{\Gamma}$ has $A(x) = 0$, and thus has $\Delta_{\text{rel}}(x) = 1$. That is, every generator of $\widehat{HF}_{\text{red}}(Y_+)$ has its relative grading increase under ϕ. Since K is nontrivial, $\widehat{\Gamma}$ has at least one vertical segment so $\widehat{HF}_{\text{red}}(Y_+)$ is nontrivial. Thus $\Sigma M_{\text{rel}}(Y_+) > \Sigma M_{\text{rel}}(Y_-)$, and $Y_+ \not\cong Y_-$. ∎

The above corollary demonstrates how powerful Proposition 4 can be; we will use this proposition to derive several more restrictions on the multicurve $\widehat{\Gamma}$. For example, Corollary 26 follows from the fact that if too many vertical segments are at height 0, the sum of all gradings will increase when ϕ is applied. On the other hand, it is easy to see that if fewer than half of all vertical segments are at height 0, then the sum of all gradings will decrease. To make this intuition more precise, recall that n_s denotes the number of vertical segments of $\widehat{\Gamma}$ at height s; by the rotational symmetry of $\widehat{\Gamma}$, $n_{-s} = n_s$. The number of generators x in $\widehat{HF}_{\text{red}}(Y_\pm)$ with $A(x) = s$ is $q \cdot n_s$. Since $k(x)$ is always nonnegative, Proposition 4 implies that $\Delta_{\text{rel}}(x) \le 1 - 2|s|$ when $A(x) = s$. It follows that if $\Sigma M_{\text{rel}}(Y_-) - \Sigma M_{\text{rel}}(Y_+) = 0$, then

$$n_0 \ge \sum_{s \neq 0} n_s(2|s| - 1) = \sum_{s=1}^{\infty} 2n_s(2s - 1) = 2n_1 + 6n_2 + 10n_3 + \cdots . \tag{9}$$

Armed with this information, we can show that large surgery slopes can never give truly cosmetic surgeries.

Theorem 27. *Let K be a nontrivial knot in S^3. If $Y_+ = S^3_{p/q}(K)$ and $Y_- = S^3_{-p/q}(K)$ are diffeomorphic and $p/q > 1$, then $p/q = 2$, $g(K) = 2$, and $n_0 = 2n_1$.*

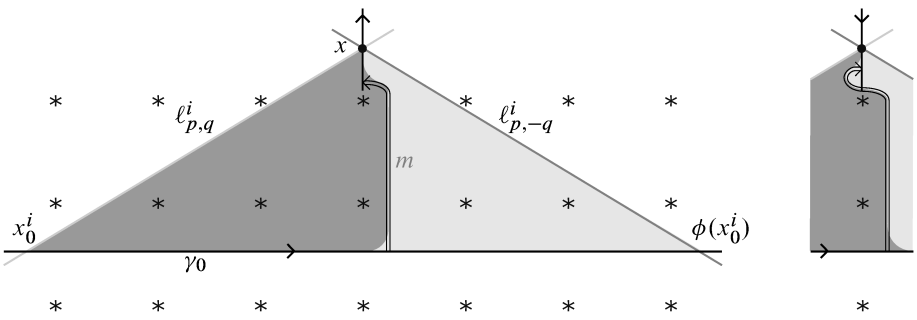

Fig. 9. Left: A computation of $\Delta_{\text{rel}}(x)$. The dark shaded bigon (which has no cusps) can be used to compute $M_{\text{rel}}(x)$, while the lightly shaded region is a cusped bigon used to compute $M_{\text{rel}}(\phi(x))$. Right: The modification to the diagram needed if the vertical segment is oriented down rather than up.

Proof. Suppose $Y_+ \cong Y_-$. In particular, there is some permutation σ on $\mathbb{Z}/p\mathbb{Z}$ such that $\widehat{HF}(Y_+, i) \simeq \widehat{HF}(Y_-, \sigma(i))$ as graded vector spaces. Note that σ can only permute spin^c structures with the same rank of $\widehat{HF}$, since

$$\widehat{HF}(Y_+, \sigma(i)) \simeq \widehat{HF}(Y_-, \sigma(i)) \simeq \widehat{HF}(Y_+, i)$$

as *ungraded* vector spaces, where the first isomorphism is given by $\phi_{\sigma(i)}$. Because $p/q > 1$, any line of slope p/q intersects any vertical segment of $\widehat{\Gamma}$ at most once. In particular, for $q \le i < p$, the line $\ell^i_{p,q}$ does not hit the vertical segments at height 0 at all, while for $0 \le j < q$ the line $\ell^j_{p,q}$ does hit the vertical segments at height 0. Thus we observe that for $j < q$ and $i \ge q$,

$$\operatorname{rk} \widehat{HF}(Y_\pm, j) \ge 1 + n_0 \quad \text{and} \quad \operatorname{rk} \widehat{HF}(Y_\pm, i) \le 1 + \sum_{s=1}^{\infty} 2n_s.$$

From (9), we find that

$$\operatorname{rk} \widehat{HF}(Y_\pm, j) - \operatorname{rk} \widehat{HF}(Y_\pm, i) \ge n_0 - \sum_{s=1}^{\infty} 2n_s \ge \sum_{s=1}^{\infty} 2n_s(2s-2) = 4n_2 + 8n_3 + \cdots.$$

This difference in dimensions is strictly positive unless $n_s = 0$ for all $s > 1$.

First suppose that $n_s > 0$ for some $s > 1$, so that the above difference is positive for any $j < q$ and $i \ge q$. The spin^c structures of Y_+ and Y_- can thus be divided by rank into two subsets, with one set having the q largest dimensions of $\widehat{HF}$ and the other set having the $p - q$ smallest dimensions of $\widehat{HF}$, and for both Y_+ and Y_-, these subsets are $\{0, \dots, q-1\}$ and $\{q, \dots, p-1\}$. The permutation σ must fix these two sets; in other words, the first q spin^c structures of Y_+ must correspond to the first q spin^c structures of Y_- under any isomorphism of $\widehat{HF}(Y_+)$ and $\widehat{HF}(Y_-)$. In particular, the sum of the d-invariants of these first q spin^c structures must agree. We have

$$\begin{aligned}\sum_{i=0}^{q-1} d(L(p,q),i) = \sum_{i=0}^{q-1} d(Y_+,i) = \sum_{i=0}^{q-1} d(Y_-,i) = \sum_{i=0}^{q-1} d(L(p,-q),i)\\ = \sum_{i=0}^{q-1} -d(L(p,q),i).\end{aligned}$$

It follows that the sum must be zero, but this is impossible by Lemma 23.

Now suppose that $n_s = 0$ for all $s > 1$. It follows that $g(K) = 2$, since the maximum s for which $n_s \ne 0$ is $g(K) - 1$, and $g(K)$ cannot be 1 by Corollary 22. Since $p/q > 1$, we have $k(x) = 0$ for any $x \in \widehat{HF}_{\mathrm{red}}(Y_+)$ (i.e. the relevant triangle does not cover any marked points). By Proposition 4,

$$\Delta_{\mathrm{rel}}(x) = \begin{cases} 1, & A(x) = 0, \\ -1, & A(x) = \pm 1. \end{cases}$$

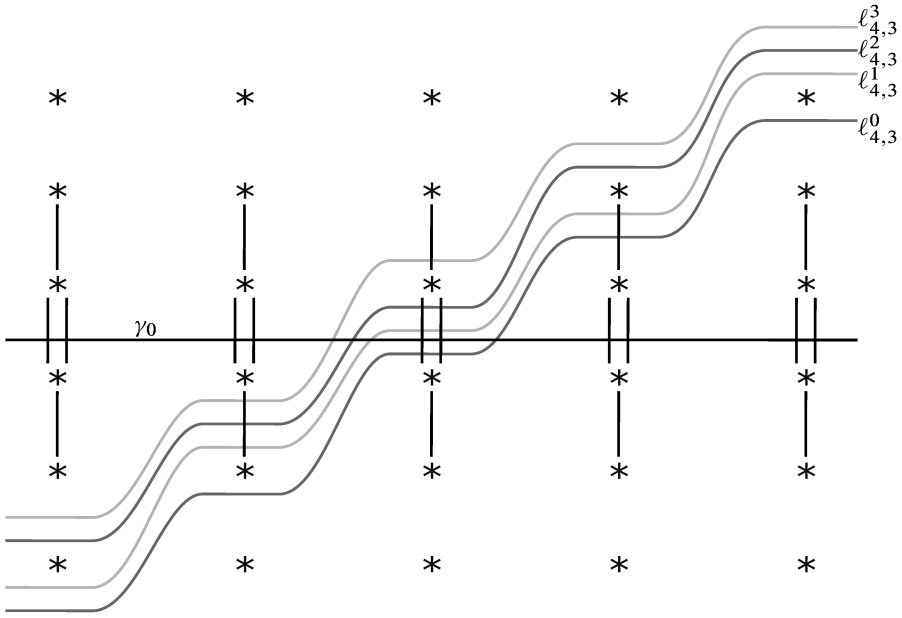

Fig. 10. Intersecting $\widehat{\Gamma}$ with four lines of slope $4/3$, assuming the knot has genus 2 with $n_0 = 2n_1$. Each vertical line in the figure represents n_1 vertical segments in $\widehat{\Gamma}$. The spin^c structures $i \in \{0, 1, 2, 3\}$ have types (b), (a), (b), and (c), respectively, as described in the proof of Theorem 27.

These grading changes must cancel when we sum over all generators of $\widehat{HF}_{\text{red}}(Y_+)$, which implies that $n_0 = n_1 + n_{-1} = 2n_1$.

There are five possible values of $\operatorname{rk} \widehat{HF}(Y_+, i)$ depending on which vertical segments $\ell^i_{p,q}$ intersects, as shown in the table below (see also Figure 10 for an example with a particular slope); this partitions the set of spin^c structures into five subsets. For each type, we can also compute the net change in relative grading, the sum of $\Delta_{\text{rel}}(x)$ over all generators x of $\widehat{HF}_{\text{red}}(Y_+, i)$.

	Heights of vertical segments hit by $\ell^i_{p,q}$	$\operatorname{rk} \widehat{HF}(Y_+, i)$	$\sum_x \Delta_{\text{rel}}(x)$
(a)	$\{1, 0, -1\}$	$1 + 4qn_1$	0
(b)	$\{1, 0\}$ or $\{0, -1\}$	$1 + 3qn_1$	qn_1
(c)	$\{1, -1\}$ or $\{0\}$	$1 + 2qn_1$	$2qn_1$ or $-2qn_1$
(d)	$\{1\}$ or $\{-1\}$	$1 + qn_1$	$-qn_1$
(e)	$\{\}$	1	0

Because σ can only permute spin^c structures with the same rank, σ must fix the subsets of spin^c structures corresponding to these five types. However, for any spin^c structure of type (b), the total relative grading strictly increases when ϕ is applied; it follows that there can be no spin^c structures of type (b). There can also be no spin^c structures of type (d) for

similar reasons. But if $p/q > 2$ there is at least one spin^c structure of type (d), namely the one defined by $i = q$, and if $1 < p/q < 2$ there is at least one spin^c structure of type (b), namely $i = 0$. Either case gives a contradiction, so p/q must be 2. ∎

Small slopes can be dealt with in a similar way. Note that for $p/q < 1$, the constant $k(x)$ in Proposition 4 is at least $s(s-1)/2$, where $s = |A(x)|$. Moreover, $k(x)$ is strictly larger than this for at least one intersection point x with $|A(x)| = s$, provided there are any such points. Since we require that $\Sigma M_{\text{rel}}(Y_-) - \Sigma M_{\text{rel}}(Y_+) = 0$, it follows that there exist constants $a_s \in \mathbb{Q}$ with $a_{-s} = a_s$ and $a_s > 2s^2 - 1$ for $s > 0$ such that

$$n_0 = \sum_{s \neq 0} n_s a_s = \sum_{s=1}^{\infty} 2n_s a_s > 2n_1 + 14n_2 + 34n_3 + \cdots . \tag{10}$$

The constants a_s could be computed exactly for any fixed p/q, but we will not need this; in fact, we will only need that $a_s > 1$ for all $s > 0$.

Theorem 28. *Let K be a nontrivial knot in S^3. If $Y_+ = S^3_{p/q}(K)$ and $Y_- = S^3_{-p/q}(K)$ are diffeomorphic and $p/q < 1$, then $p = 1$.*

Proof. Suppose to the contrary that $Y_+ \cong Y_-$ and $p \neq 1$. Let $q = mp + r$ with $0 \leq r < p$; since p and q are relatively prime, $r > 0$. Any line of slope p/q hits any vertical segment in $\widehat{\Gamma}$ either m or $m+1$ times; let $c_s^i \in \{m, m+1\}$ denote the number of times the line $\ell^i_{p,\pm q}$ intersects a vertical segment at height s. We have

$$\operatorname{rk} \widehat{HF}(Y_\pm, i) = 1 + \sum_{s=-\infty}^{\infty} c_s^i n_s.$$

If $0 \leq i < r$, it is easy to see that $c_0^i = m+1$, while if $r \leq j < p$ then $c_0^j = m$. Since $c_s^i - c_s^j \geq -1$ for any s,

$$\operatorname{rk} \widehat{HF}(Y_\pm, i) - \operatorname{rk} \widehat{HF}(Y_\pm, j) = \sum_{s=-\infty}^{\infty} (c_s^i - c_s^j) n_s \geq n_0 - \sum_{s=1}^{\infty} 2n_s > 0,$$

where the last inequality uses (10). In other words, the first r spin^c structures have ranks strictly bigger than each of the remaining $p - r$ spin^c structures. It follows that the permutation σ corresponding to the reindexing of spin^c structures under any isomorphism from $\widehat{HF}(Y_+)$ to $\widehat{HF}(Y_-)$ must preserve the first r spin^c structures as a set. In particular,

$$\sum_{i=0}^{r-1} d(Y_+, i) = \sum_{i=0}^{r-1} d(Y_-, i).$$

As in the proof of Theorem 27, this implies that

$$0 = \sum_{i=0}^{r-1} d(L(p,q), i) = \sum_{i=0}^{r-1} d(L(p,r), i).$$

But $r \equiv q \equiv -1 \pmod{p}$, so this is impossible by Lemma 23. ∎

When we restrict to $p = 1$ we can compute explicit formulas for the net change in relative grading under ϕ, and this determines q exactly (for a given knot) if $\pm 1/q$ is a pair of truly cosmetic surgery slopes.

Proposition 29. *Suppose $S^3_{1/q}(K) \cong S^3_{-1/q}(K)$. As above, let n_s be the number of vertical segments in $\widehat{\Gamma}$ at height s. Then*

$$q = \frac{n_0 + 2\sum_{s=1}^{\infty} n_s}{4\sum_{s=1}^{\infty} s^2 n_s}.$$

Proof. This is a straightforward consequence of Proposition 4 and the fact that $\sum_x \Delta_{\text{rel}}(x) = 0$. Note that for slope $1/q$ there are q intersections of line $\ell_{1,q}$ with any vertical segment at height $s > 0$, and if we label these intersections by $i = 0, \ldots, q-1$, the constant $k(x)$ for the ith intersection is

$$[q(s-1) + q(s-2) + \cdots + q] + is = \frac{qs(s-1)}{2} + is,$$

and the sum over all q of these points of $4k(x) + 2A(x) - 1$ is

$$4\left[\frac{q^2 s(s-1)}{2} + \frac{sq(q-1)}{2}\right] + 2qs - q = 2q^2 s^2 - q.$$

By symmetry the contribution to $\sum_x \Delta_{\text{rel}}(x)$ of a vertical segment at height $-s$ is the same as that of a vertical segment at height s. Each vertical segment at height 0 contributes q intersection points, each with $\Delta_{\text{rel}}(x) = 1$. The condition that $\sum_x \Delta_{\text{rel}}(x) = 0$ can now be stated as

$$qn_0 - 2\sum_{s=1}^{\infty}(2q^2 s^2 - q)n_s = 0;$$

solving this equation for q gives the desired result. ∎

Note that since $n_s = 0$ for $|s| \geq g(K)$, the infinite sums above can be truncated for any particular example.

Another powerful consequence of Proposition 4 is a bound on q and the genus of K. Theorems 27 and 28 and Proposition 29 rely on the fact that the sums of relative gradings $\Sigma M_{\text{rel}}(Y_+)$ and $\Sigma M_{\text{rel}}(Y_-)$ should agree, and thus for every generator x of $\widehat{HF}_{\text{red}}(Y_+)$ with $\Delta_{\text{rel}}(x) = -n < 0$, there must be n generators $y_1, \ldots, y_n$ with $\Delta_{\text{rel}}(y_i) = 1$. But the set of relative gradings is an invariant, not just its sum, so in fact $M_{\text{rel}}(Y_+) = M_{\text{rel}}(Y_-)$ as multisets. This lets us say more:

Lemma 30. *Suppose $Y_+ \cong Y_-$. If $\widehat{HF}_{\text{red}}(Y_+)$ contains a generator x with $\Delta_{\text{rel}}(x) = -n < 0$ and $M_{\text{rel}}(x) = m$, then it must contain generators $y_1, \ldots, y_n$ with $\Delta_{\text{rel}}(y_i) = 1$ and $M_{\text{rel}}(y_i) = m - i$.*

Proof. If $\psi : \widehat{HF}(Y_+) \to \widehat{HF}(Y_-)$ is a grading preserving isomorphism, then $\psi^{-1} \circ \phi$ determines a permutation on the multiset of gradings $M_{\text{rel}}(Y_+)$. This permutation takes an element m of $M_{\text{rel}}(Y_+)$ (corresponding to the grading of the generator x) to $m - n$.

The cycle containing this element must eventually return to m, but at each step the relative grading can increase by at most one by Proposition 4. Thus for each $m-i$ between $m-n$ and $m-1$, $m-i$ appears in the cycle followed by $m-i+1$; each such element $m-i$ corresponds to M_{rel} for some generator with positive Δ_{rel}, and we take this generator to be y_i. ∎

We now need to relate M_{rel} of generators of $\widehat{HF}(Y_+)$ with $\Delta_{\text{rel}}=1$ to the knot Floer homology of K. Recall that the δ-grading refers to the difference between the Alexander and Maslov gradings.

Lemma 31. *If x is a generator of $\widehat{HF}(Y_+)$ coming from a vertical segment of height 0 and $M_{\text{rel}}(x)=m$, then there is an intersection point in $\widehat{\Gamma}\cap\mu$ (i.e. a generator of $\widehat{HFK}(K)$) with δ-grading at least $-m-1$, and there is one with δ-grading at most $-m$.*

Proof. Let V be the vertical segment in $\widehat{\Gamma}$ containing x. We may assume there is a grading arrow in $\widehat{\Gamma}$ moving upward from γ_0 to V, as shown in Figure 11 (a) if m is odd and 11 (d) if m is even. Consider first the case that m is odd, with $m=2k-1$; then the arrow has weight k and approaches V from the right. We will assume that $\widehat{\Gamma}$ is in minimal position with μ and note that V lies either to the right of μ, as in Figure 11 (b), or to the left of μ, as in Figure 11 (c). In either case, let x' and x'' in $\widehat{\Gamma}\cap\mu$ be the first intersections with μ when following $\widehat{\Gamma}$ upward and downward, respectively, from V. Note that $A(x')\geq 0$ and $A(x'')\leq 0$.

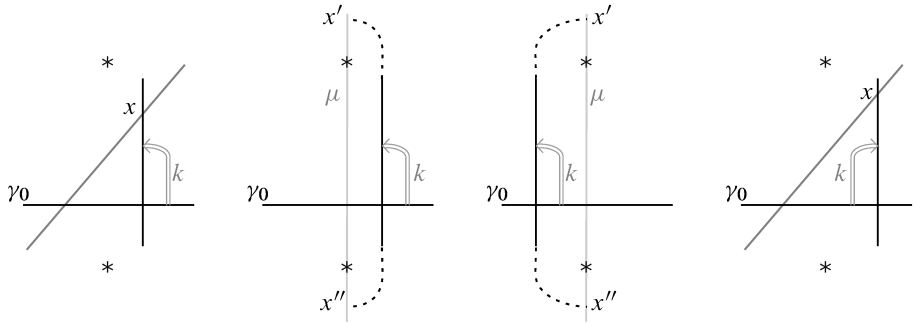

Fig. 11. (a) A grading arrow such that $M_{\text{rel}}(x)=m=2k-1$; (b, c) x' and x'', the first intersections with μ above and below the vertical segment containing x; (d) a grading arrow such that $M_{\text{rel}}(x)=m=2k$.

Let x_0 denote the intersection of γ_0 with μ, which by definition has Maslov grading $M(x_0)=0$. If V is to the right of μ, there is a bigon from x_0 to x' whose boundary traverses the grading arrow and which passes $A(x')$ marked points along μ; from this and (3) in Section 2.4 we compute

$$M(x')=M(x_0)+2k+2A(x')-1=m+2A(x'),$$
$$\delta(x')=A(x')-M(x')=-m-A(x')\leq -m.$$

The bigon from x'' to x', which passes $A(x') - A(x'')$ marked points along μ, gives

$$M(x'') = M(x') + 1 - 2(A(x') - A(x'')) = m + 2A(x'') + 1,$$
$$\delta(x'') = A(x'') - M(x'') = -m - 1 - A(x'') \geq -m - 1.$$

If instead V lies to the left of μ then a bigon from x_0 to x'' gives

$$M(x'') = 2k - 1 = m, \quad \delta(x'') = A(x'') - m \leq -m,$$

and a bigon from x' to x'' gives

$$M(x') = M(x'') + 1 = m + 1, \quad \delta(x') = A(x') - m - 1 \geq -m - 1.$$

The case that $m = 2k$ is similar. The grading arrow must approach V from the left, as in Figure 11 (d), which adds a cusp to the boundary of the bigons considered above which run over the grading arrow. This increases $M(x')$ and $M(x'')$ by 1 (computing in terms of k), which decreases the δ-gradings by 1, but m is also increased from $2k - 1$ to $2k$, so the conclusion still holds. ∎

We are now ready to prove the promised bounds on q and $g(K)$. Recall that $\mathrm{th}(K)$ denotes the Heegaard Floer thickness of K, that is,

$$\mathrm{th}(K) = \max\{\delta(x) \mid x \in \widehat{HFK}(K)\} - \min\{\delta(x) \mid x \in \widehat{HFK}(K)\}.$$

Theorem 32. *Let K be a nontrivial knot in S^3 of genus g. If $Y_+ = S^3_{1/q}(K)$ and $Y_- = S^3_{-1/q}(K)$ are diffeomorphic, then*

$$\mathrm{th}(K) \geq 2qg(g-1) - 2g.$$

Proof. The maximum height attained by $\widehat{\Gamma}$ is g so there are at least two vertical segments at height $g - 1$, one on either side of a maximum of $\widehat{\Gamma}$. Consider the highest intersection point of each of these two vertical segments with any $\ell^i_{1,q}$; call these x and x' as shown in Figure 12. By construction, $A(x) = A(x') = g - 1$. Counting the marked points in the closure of the triangle in Figure 12 and removing those on the boundary gives

$$k(x) = k(x') = [q(g-1) + q(g-2) + \cdots + q(1)] - (g-1) = q\frac{g(g-1)}{2} - (g-1).$$

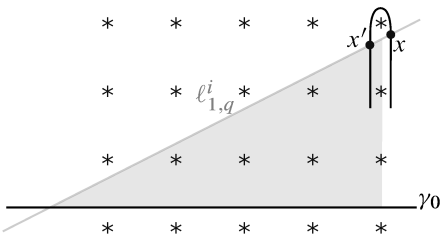

Fig. 12. Generators x and x' with minimal Δ_{rel} in $\widehat{HF}_{\mathrm{red}}(Y_+)$ for surgery slope $1/q$. Here $q = 2$ and $g(K) = 3$.

Let d denote $-\Delta_{\mathrm{rel}}(x) = -\Delta_{\mathrm{rel}}(x')$. We have

$$\begin{aligned} d &= 4k(x) + 2A(x) - 1 = 4\left(\frac{qg(g-1)}{2} - (g-1)\right) + 2(g-1) - 1 \\ &= 2qg(g-1) - 2g + 1. \end{aligned}$$

Let m be $M_{\mathrm{rel}}(x)$; the small bigon from x to x' covering one puncture tells us that $M_{\mathrm{rel}}(x') = M_{\mathrm{rel}}(x) + 1 = m + 1$. By Lemma 30, there are generators $y_0, \dots, y_d$ of $\widehat{HF}_{\mathrm{red}}(Y_+)$ with $M_{\mathrm{rel}}(y_i) = m - i$ and $\Delta_{\mathrm{rel}}(y_i) = +1$. Applying Lemma 31 to y_0 and y_d, we find that $\widehat{HFK}(K)$ contains generators z_1 and z_2 with $\delta(z_1) \le -m$ and $\delta(z_2) \ge -m - 1 + d$. Thus $\mathrm{th}(K) \ge d - 1$. ∎

Theorems 27 and 28, Proposition 29, and Theorem 32 combine to give Theorem 2 in the introduction.

5. Explicit obstructions in terms of $\widehat{\Gamma}$

5.1. General constraints on $\widehat{\Gamma}$

In addition to Theorem 2, it is helpful to have explicit conditions on a knot K, in terms of its knot Floer invariant, that ensure that K admits no truly cosmetic surgeries at all. Several such conditions are already implicit in what has been discussed so far. One condition comes from Theorem 21, namely that the curve γ_0 is horizontal. Another condition follows from the bounds in Theorem 2: if $g = g(K) \neq 2$ and $\mathrm{th}(K) + 2g < 2g(g-1)$, then K admits no truly cosmetic surgeries. We have also seen that Proposition 4 and the fact that $\sum_x \Delta_{\mathrm{rel}}(x) = 0$ places constraints on the numbers n_s of vertical segments in $\widehat{\Gamma}(K)$ at height s, including the inequality (9). In particular, if fewer than half of all vertical segments occur at height 0, then K cannot admit any truly cosmetic surgeries. For the slopes $1/q$ the inequality (9) can be improved to an equation that can be solved for q, giving rise to Proposition 29. This places a further constraint on the existence of truly cosmetic surgeries which is implicit in Theorem 2: if the quantity

$$\frac{n_0 + 2\sum_{s=1}^{\infty} n_s}{4\sum_{s=1}^{\infty} s^2 n_s}$$

is not a positive integer for a given knot K, then K does not admit truly cosmetic surgeries.

To arrive at Proposition 29 and the other constraints mentioned above, we only assumed that $S^3_r(K)$ and $S^3_{-r}(K)$ have the same *sum* of all relative gradings. By considering the *set* of all relative gradings, we could impose further constraints on K. For example, for $r = 2$, it is not enough to have two vertical segments at height 0 for each vertical segment at height 1, we also require that vertical segments at height 0 give rise to generators in $S^3_r(K)$ that have grading 1 less than those coming from the vertical segment at height 1. Unfortunately, it is difficult to state such conditions purely in terms of the knot Floer homology of K; this is partly because there is not a perfect correspondence between

gradings of elements of the surgery coming from a vertical segment and gradings in knot Floer homology.

We will see that this difficulty can be overcome in a particular special case. We remark that while it might be possible to state some additional (likely messy) conditions on $\widehat{\Gamma}(K)$ in the general case, it is not worth doing so. In practice, to achieve this fine an obstruction on K it is easiest to simply compute absolutely graded HF^+ for each of the finitely many surgery pairs allowed by Theorem 2 and check if they agree for any pair. In this way, we can extract the maximum information from Heegaard Floer homology: for any knot we can always either rule out cosmetic surgeries or conclude that Heegaard Floer homology cannot rule out some potential pair. This can be thought of as a condition on the knot Floer homology of K (by the surgery formula, HF^+ of the relevant surgeries is determined by knot Floer homology), though it is a condition that requires some computation to check.

5.2. *Further constraints for simple figure eight curves*

There is one situation where it is convenient to state additional constraints purely in terms of knot Floer homology, and that is when the underlying curve set for $\widehat{\Gamma}$ consists only of γ_0 and simple figure eight curves. As mentioned in Remark 9, this is common in practice.

A simple figure eight component γ_i of $\boldsymbol{\gamma}$ intersects μ four times, corresponding to four generators of the knot Floer homology of K. These generators all have the same δ-grading, so it makes sense to talk about the δ-grading of the curve γ_i. The height of a simple figure eight component is the height at which it is centered, which is the Alexander grading of two of the four generators. Let e_s^d denote the number of simple figure eight components in $\boldsymbol{\gamma}$ at height s with δ-grading d, and let $e_s = \sum_{d\in\mathbb{Z}} e_s^d$ be the total number of simple figure eights at height s. Each simple figure eight curve contributes two vertical segments at height s, so if we assume γ_0 is horizontal and all other γ_i's are simple figure eights then $n_s = 2e_s$. We will assume the self-intersection in a simple figure eight curve γ_i occurs below the vertical segments, as in Figure 13, and with this understanding we will refer to the *left* and *right* vertical segments coming from γ_i. The relative grading of γ_i is determined by its δ-grading. In particular, if γ_i is a simple figure eight at height $s \geq 0$ with δ-grading d, we can add a consistent grading arrow from γ_0 to one of the vertical segments in γ_i as shown in Figure 13 (a, b); the arrow passes to the right of any marked points up to height s and ends on the right vertical segment and carries the weight $\frac{-d-s}{2}$ if $d+s$ is even, or it ends on the left vertical segment and carries the weight $\frac{1-d-s}{2}$ if $d+s$ is odd. The case of height $s<0$ is similar, except that the arrow stays to the left of the marked points and s is replaced with $|s|$ in the arrow weights.

Remark 33. This gives an alternative (and much simpler) way of encoding grading information in $\widehat{\Gamma}(K)$ when all curves other than γ_0 are simple figure eights: instead of decorating the set of curves with a collection of grading arrows, we can simply decorate each curve other than γ_0 with an integer, its δ-grading.

We now relate the δ-grading d of a simple figure eight component γ_i at height s to the relative grading of generators of $\widehat{HF}(Y_+)$ coming from γ_i. Note that for each generator x

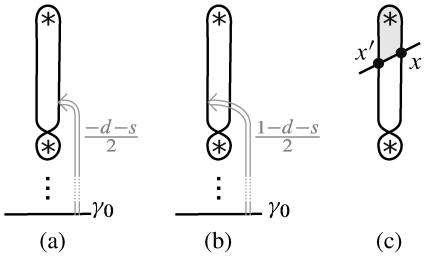

Fig. 13. (a) A grading arrow for a simple figure eight component at height s with δ-grading d if $d+s$ is even; (b) a grading arrow if $d+s$ is odd; (c) a pair x, x' of intersection points representing generators of $\widehat{HF}(Y_+)$, with $M_{\text{rel}}(x') = M_{\text{rel}}(x) + 1$.

coming from the right vertical segment of γ_i, there is a corresponding generator x' coming from the left vertical segment, as shown in Figure 13 (c). These intersection points are connected by a small bigon covering one marked point, so that $M_{\text{rel}}(x') = M_{\text{rel}}(x) + 1$. It is clear that $k(x) = k(x')$ and $A(x) = A(x') = s$, and thus $\Delta_{\text{rel}}(x) = \Delta_{\text{rel}}(x') = 1 - 4k(x) - 2s$. We calculate that $M_{\text{rel}}(x) = -1 + 2k(x) + s - d$, regardless of the parity of $s + d$. For any given slope r, it is possible to compute the $k(x)$ for all intersection points of lines of slope r and state an obstruction to $S^3_r(K)$ and $S^3_{-r}(K)$ agreeing purely in terms of the quantities e^d_s for K; we will only do this for the slope $r = 1$, where every intersection point x coming from a simple figure eight at height s has $k(x) = s(s-1)/2$.

Proposition 34. *Suppose that $\widehat{\Gamma}(K)$ consists only of γ_0 and simple figure eight curves, and let e^d_s be the number of simple figure eight components with height s and δ-grading d. If $S^3_{+1}(K) \cong S^3_{-1}(K)$, then γ_0 is horizontal and for every $D \in \mathbb{Z}$, we must have*

$$e^D_0 = \sum_{s \neq 0} \sum_{d=-D-s^2+1}^{-D+s^2-1} e^d_s .$$

Proof. We already know that γ_0 is horizontal by Theorem 21. We now consider the graded contribution of the simple figure eight components to each surgery. It is enough to consider only one generator of $\widehat{HF}_{\text{red}}(S^3_{+1}(K))$ coming from each simple figure eight component, the one coming from the right vertical segment, since the generators coming from the left vertical segments behave exactly the same with the grading shifted up by 1. With this in mind, a simple figure eight component with height s and δ-grading d contributes a generator x with $M_{\text{rel}}(x) = s^2 - 1 - d$ and $\Delta_{\text{rel}}(x) = 1 - 2s^2$. By Lemma 30, this must be counteracted by a chain of generators coming from height 0 vertical segments with gradings ranging from $-d - s^2$ up to $-d + s^2 - 2$. These must come from a chain of height 0 simple figure eight components with δ-gradings $-d - s^2 + 1, -d - s^2 + 2, \ldots, -d + s^2 - 1$ (we are still considering only the right vertical segment of each figure eight). All height 0 figure eight components must be accounted for in one of these chains, and we see that there is a contribution to e^D_0 for each figure eight component with height $s \neq 0$ and grading d with $d \geq -D - s^2 + 1$ and $d \leq -D + s^2 - 1$. ∎

Restricting to knots with small thickness gives the following condition:

Proposition 35. *Suppose, as in the previous proposition, that* $\widehat{\Gamma}(K)$ *consists only of* γ_0 *and simple figure eight components, and suppose* $\mathrm{th}(K) < 4$. *If* K *admits any truly cosmetic surgery, then* $e_s^d = 0$ *for all* $|s| > 1$, *and* $e_0^d = e_1^d + e_{-1}^d = 2e_1^d$ *for every* δ*-grading* d.

Proof. By Theorem 2, we know that $g(K) = 2$, which implies $e_s^d = 0$ if $|s| \geq 2$. We also know that the only possible truly cosmetic surgeries have slopes ± 1 or ± 2. In the first case, we apply Proposition 34 to get the desired result. For the latter case, note that intersections x between lines of slope 2 and vertical segments at height 1 or -1 still have $k(x) = 0$, and so $M_{\mathrm{rel}}(x)$ and $\Delta_{\mathrm{rel}}(x)$ are the same as in the slope 1 case, and the reasoning in the proof of Proposition 34 applies. ∎

There is a sort of Heegaard Floer converse to this statement: if $g(K) = 2$, and $e_0^d = 2e_1^d$ for all gradings d, then Heegaard Floer homology can not distinguish either pair—that is, $\widehat{HF}(S_r^3(K))$ and $\widehat{HF}(S_{-r}^3(K))$ are isomorphic as absolutely graded vector spaces for $r \in \{1, 2\}$. We compute this explicitly for the example of 9_{44} below, which serves as a model computation for the general case. One might hope that upgrading to HF^+ would help in this situation, but it does not. In fact, when γ_0 is horizontal and all other curves in $\widehat{\Gamma}(K)$ are simple figure eight components, HF^+ is determined by $\widehat{HF}$ for any surgery on K.

In the case of thin knots this condition can be given purely in terms of the Alexander polynomial, as stated in the introduction.

Proof of Theorem 5. We use Proposition 35, though since K is thin there is only one occupied δ-grading, which must be 0 if γ_0 is horizontal. For K to admit truly cosmetic surgeries, we must have some number n of simple figure eights at height 1 (and at height -1, by symmetry), and $2n$ simple figure eight components at height 0. It is easy to compute $\Delta_K(t)$ from this information and see that it has the desired form. (Conversely, for thin knots $\widehat{\Gamma}(K)$ is determined by $\Delta_K(t)$ and $\sigma(K)$; for $\sigma(K) = 0$ and $\Delta_K(t)$ as in the conclusion of the theorem, it is easy to check the γ_0 is horizontal, $e_1^0 = e_{-1}^0 = n$, and $e_0^0 = 2n$, and thus Heegaard Floer homology does not distinguish ± 1 surgeries or ± 2 surgeries.) ∎

5.3. Unobstructed knots

We conclude this section by demonstrating that Theorem 2 cannot be substantially improved using Heegaard Floer homology alone. We first note that there exist knots for which Heegaard Floer homology does not distinguish ± 1 surgeries or ± 2 surgeries. For example, consider the knot 9_{44} shown in Figure 14 (a). This example appeared in [19], where it was first observed that $HF^+(S_1^3(9_{44})) \cong HF^+(S_{-1}^3(9_{44}))$. It turns out that this example is representative of all currently known examples for which Heegaard Floer homology does not obstruct truly cosmetic surgeries, so we now examine this example

in detail. The knot Floer invariant $\widehat{\Gamma}(9_{44})$ is shown in Figure 14 (b) (see also Figure 3); γ_0 is horizontal and there are four simple figure eight components, with $e^0_1 = e^0_{-1} = 1$ and $e^0_0 = 2$.

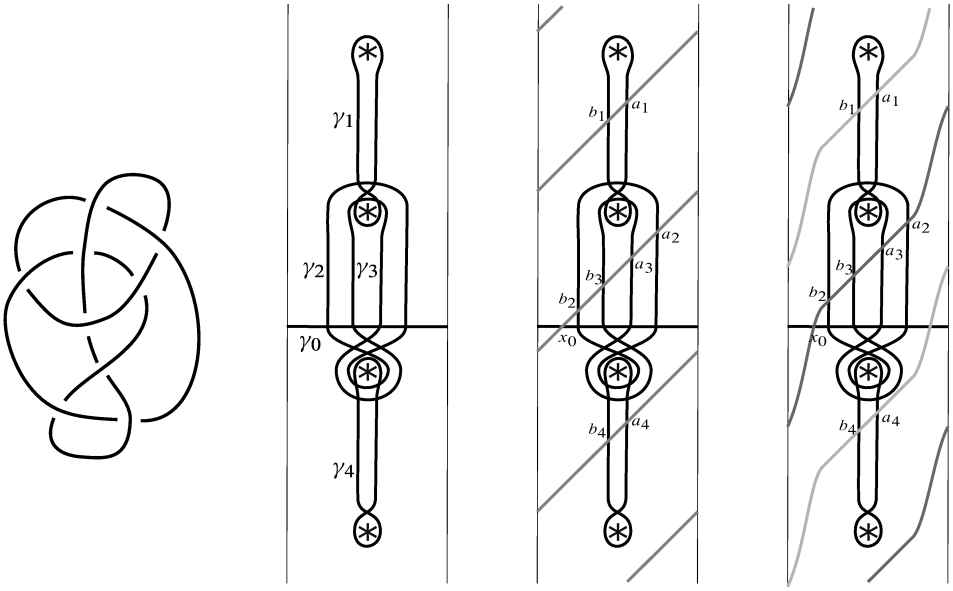

Fig. 14. (a) The knot 9_{44}; (b) the invariant $\widehat{\Gamma}(9_{44})$ (all simple figure eight components have δ-grading 0); (c) the computation of $\widehat{HF}(S^3_1(9_{44}))$; (d) the computation of $\widehat{HF}(S^3_2(9_{44}))$.

Let $Y_\pm$ denote $S^3_{\pm 1}(9_{44})$; we compute $\widehat{HF}(Y_+)$ using part (c) of Figure 14, which shows the intersection of $\widehat{\Gamma}(9_{44})$ with a line of slope $+1$. It is easy to see that, in addition to the generator x_0 coming from γ_0, $\widehat{HF}(Y_+)$ has a pair of generators (a_i, b_i) coming from each of the simple figure eight curves, with $M_{\text{rel}}(b_i) = M_{\text{rel}}(a_i) + 1$. For the two extremal figure eights (that is, for $i \in \{1, 4\}$), we have $M_{\text{rel}}(a_i) = 0$, $M_{\text{rel}}(b_i) = 1$, and $\Delta_{\text{rel}}(a_i) = \Delta_{\text{rel}}(b_i) = -1$. For the two figure eights at height 0 (that is, for $i \in \{2, 3\}$), we have $M_{\text{rel}}(a_i) = -1$, $M_{\text{rel}}(b_i) = 0$, and $\Delta_{\text{rel}}(a_i) = \Delta_{\text{rel}}(b_i) = 1$. Applying ϕ with the given grading changes, we see that $\widehat{HF}(Y_-)$ has $M_{\text{rel}}(\phi(x_0)) = 0$, $M_{\text{rel}}(\phi(a_i)) = -1$ for $i \in \{1, 4\}$, $M_{\text{rel}}(\phi(b_i)) = 0$ for $i \in \{1, 4\}$, $M_{\text{rel}}(\phi(a_i)) = 0$ for $i \in \{2, 3\}$, and $M_{\text{rel}}(\phi(b_i)) = 1$ for $i \in \{2, 3\}$. Thus $\widehat{HF}(Y_+)$ and $\widehat{HF}(Y_-)$ agree as relatively graded vector spaces. The absolute grading agrees as well since $d(Y_+) = 0 = d(Y_-)$. Note that the computation is similar for any genus 2 knot for which $\widehat{\Gamma}$ contains only a horizontal γ_0 and simple figure eight curves and for which $e^d_0 = 2e^d_1 = 2e^d_{-1}$ for each δ-grading d: each height 1 or height -1 figure eight with grading d contributes a pair (a_i, b_i) with $M_{\text{rel}}(a_i) = -d$, $M_{\text{rel}}(b_i) = 1 - d$, and $\Delta_{\text{rel}}(a_i) = \Delta_{\text{rel}}(b_i) = -1$, and each height 0 figure eight with grading d contributes a pair (a_i, b_i) with $M_{\text{rel}}(a_i) = -1 - d$, $M_{\text{rel}}(b_i) = -d$, and $\Delta_{\text{rel}}(a_i) = \Delta_{\text{rel}}(b_i) = +1$.

If instead we consider $Y_\pm = S^3_{\pm 2}(9_{44})$, the computation is almost identical except that the generators now split into two spin^c structures. We compute $\widehat{HF}(Y_+)$ using part (d) of Figure 14, which shows the intersection of $\widehat{\Gamma}(9_{44})$ with two lines of slope $+2$ (one for each spin^c structure). Ignoring the generators coming from γ_0, we have $\widehat{HF}_{\text{red}}(Y_+, 0)$ generated by the four intersections with height 0 figure eights (labeled b_2, b_3, a_3, and a_2 in Figure 14), which have relative gradings $\{0, 0, -1, -1\}$ and Δ_{rel} of 1, while $\widehat{HF}_{\text{red}}(Y_+, 1)$ is generated by the four intersections with the extremal figure eights (labeled b_1, b_4, a_1, and a_4 in Figure 14), which have relative gradings $\{1, 1, 0, 0\}$ and Δ_{rel} of -1. Thus $\widehat{HF}(Y_+, 0)$ and $\widehat{HF}(Y_-, 1)$ agree as relatively graded vector spaces, as do $\widehat{HF}(Y_+, 1)$ and $\widehat{HF}(Y_-, 0)$; in other words, any graded isomorphism from $\widehat{HF}(Y_+)$ to $\widehat{HF}(Y_+)$ must permute the two spin^c structures. This could potentially conflict with the absolute grading, but it does not, since

$$d(Y_+, 0) = d(L(2, 1), 0) = \tfrac{1}{4} = d(L(2, -1), 1) = d(Y_-, 1),$$
$$d(Y_+, 1) = d(L(2, 1), 1) = -\tfrac{1}{4} = d(L(2, -1), 0) = d(Y_-, 0).$$

Once again, the computation is essentially the same for any genus 2 knot for which $\widehat{\Gamma}$ contains only a horizontal γ_0 and simple figure eight curves and for which $e_0^d = 2e_1^d = 2e_{-1}^d$ for each δ-grading d. There are other knots which satisfy this property; Table 1 gives 337 such knots. Surprisingly, these are the only knots the author is currently aware of for which Heegaard Floer homology does not obstruct all truly cosmetic surgeries. In particular, an example for which $\widehat{HF}(S^3_{1/q}(K))$ agrees with $\widehat{HF}(S^3_{-1/q}(K))$ for $q > 1$ or for a knot with $g(K) > 2$ has not yet been found. We note that it is possible to construct a decorated immersed curve $\widehat{\Gamma}$ which would allow for truly cosmetic surgeries with $g > 2$ or $q > 1$, but it is not known whether such curves occur as the invariant for a knot in S^3. To construct such a curve for some $g \geq 2$ and $q \geq 1$, we can place q simple figure eights at height $g - 1$ and, for symmetry, another q simple eights at height $1 - g$. Each of these figure eights produces q pairs of points in the intersection with $\ell_{1,q}$. Indexing these pairs of points by $1 \leq i \leq q$, each pair has some relative gradings m_i and $m_i + 1$ and some grading shift $\Delta_{\text{rel}} = -\Delta_i < 0$. For each i, we then add $2\Delta_i$ simple figure eights at height 0, with two in each δ-grading from $-m_i$ to $-m_i + \Delta_i - 1$. Each of these figure eights at a δ-grading d produces a pair of intersection points with relative gradings $-d - 1$ and $-d$ and with $\Delta_{\text{rel}} = +1$; it is straightforward to check that these grading increases counteract the grading decreases from the extremal figure eights when ϕ is applied. Thus for the resulting graded multicurve $\widehat{\Gamma}$, the Floer homologies $HF(\widehat{\Gamma}, \ell_{1,q})$ and $HF(\widehat{\Gamma}, \ell_{1,-q})$ agree as graded vector spaces.

6. Computational results

One important consequence of Theorem 2 is that for any given knot K, cosmetic surgeries on K are ruled out for all but possibly a finite number of pairs of slopes. In practice, cosmetic surgeries are obstructed outright for the vast majority of knots, and for the remaining

knots the "finite number" of possible pairs that need to be checked is quite small, often just 2. Thus checking the cosmetic surgery conjecture on any finite set of knots reduces to distingushing a small number of pairs of manifolds. Computing, say, hyperbolic invariants for these remaining pairs is very tractable, and often this is sufficient to rule out the remaining surgeries. To demonstrate this, we check the following:

Theorem 36. *The cosmetic surgery conjecture holds for all prime knots with at most* 16 *crossings.*

Proof. We computed the $UV = 0$ knot Floer complex for all ≤ 16 crossing prime knots, using a program of Szabó, and then checked each against the obstructions described in this paper.[2] Recall that the $UV = 0$ knot Floer complex of K is equivalent to the immersed curve invariant $\widehat{\Gamma}(K)$. We make two observations from these computations:

- The maximum thickness of any prime knot up to 16 crossings is 2.
- For each genus 2 knot up to 16 crossings, $\widehat{\Gamma}(K)$ contains only simple figure eight components besides γ_0.

The first observation tells us immediately that we only need to consider genus 2 knots and we only need to consider the slopes ± 1 and ± 2. The second observation tells us that for these knots we can use the obstruction in Proposition 35.

There are 1,701,935 knots up to 16 crossings. We note that the results of Ni and Wu (specifically conclusion (i) in Theorem 20) already verify the conjecture for over two-thirds of these knots: after restricting to knots with $\tau(K) = 0$, we are left with 449,417 knots (requiring that $\epsilon = 0$ rather than $\tau = 0$ eliminates a further 38 knots, leaving 449,379). It turns out that the obstructions coming from Theorem 2 are much stronger. Among knots with $\epsilon(K) = 0$, requiring also that $g(K) = 2$ reduces the list to 3,316. Finally, the obstruction in Proposition 35 rules out truly cosmetic surgery on all but 337 of these knots. The remaining knots are listed in Table 1. Thus we have reduced to 674 possible pairs of cosmetic surgeries, ± 1 and ± 2 surgeries on each of these 337 knots.

This is the best that Heegaard Floer techniques alone can tell us; as noted in the previous section, for any knot K satisfying the constraint in Proposition 35, Heegaard Floer homology cannot distinguish $S^3_{+1}(K)$ from $S^3_{-1}(K)$, nor can it distinguish $S^3_{+2}(K)$ from $S^3_{-2}(K)$. So these last examples must be ruled out using other methods. Computing the hyperbolic volume for the manifolds in question using SnapPy, we find that this distinguishes every pair except for the surgeries on four knots: 10_{33}, $16n600112$, $16n786382$, and $16n988939$. These knots are amphichiral, so $+r$ surgery and $-r$ surgery can never be distinguished by hyperbolic volume. For these manifolds, the ± 1 and ± 2 surgery pairs on each of these four knots are distinguished by the Chern–Simons invariant, also computed by SnapPy.[3] ∎

[2] All code used for these computations is available at https://github.com/hanselman/CFK-immersed-curves.

[3] The author thanks Dave Futer for suggesting the use of the Chern–Simons invariant for amphichiral examples.

Remark 37. We caution that Theorem 36 depends on computer calculations, some of which are nonverified. The computation of the knot Floer complex is combinatorial in nature, but SnapPy uses numerical methods to compute hyperbolic invariants so these computations should not be taken as rigorous proof. SnapPy does offer verified computation, via interval arithmetic, of hyperbolic volume but not of the Chern–Simons invariant. Thus, to be fully rigorous, the knots 10_{33}, $16n600112$, $16n786382$, and $16n988939$ should be excluded from Theorem 36. We do note that SnapPy *estimates* the Chern–Simons computations are accurate to 10 decimal places, and for each relevant pair of manifolds the values differ by at least .001.

Remark 38. Some of the 337 knots listed in Table 1 can be ruled out using the Jones polynomial as in [13], instead of using hyperbolic volumes. Unfortunately, this does not help with any of the four knots for which the Chern–Simons invariant was needed.

The result above considers prime knots, but Theorem 2 is also very good at obstructing truly cosmetic surgeries on connected sums. In fact, with only a little more work, we can rule out cosmetic surgeries on all knots whose prime summands have at most 16 crossings. The following is equivalent to Theorem 6 stated in the introduction.

Theorem 39. *The cosmetic surgery conjecture holds for the connected sum of any number of prime knots each with* ≤ 16 *crossings.*

Proof. Suppose that K has $n > 1$ prime summands, each with at most 16 crossings. It follows that $g(K) \geq n$ and $\mathrm{th}(K) \leq 2n$, since both genus and thickness are additive with respect to connected sum and the maximum thickness for knots up to 16 crossings is 2. Suppose $\pm r$ is a pair of truly cosmetic surgery slopes for K. If $g(K) > 2$, then by Theorem 2 we must have $r = 1/q$ with

$$q \leq \frac{\mathrm{th}(K)}{2g(K)(g(K)-1)} + \frac{1}{g(K)-1} \leq \frac{2n}{2n(n-1)} + \frac{1}{n-1} = \frac{2}{n-1}.$$

Since q must be ≥ 1, it follows that $n \leq 3$.

Moreover, if $n = 3$ then $g(K) = 3$ and $q = 1$. If $n = 2$ then $g(K) = 2$ but q can be 1 or 2; we must also consider the case $g(K) = 2$ and $r = 2$. By Proposition 29, if $n = 3$ we require that $e_0 = 2e_1 + 7e_2$, while if $n = 2$ we require $e_0 = 2e_1$ if $r \in \{1, 2\}$ and $e_0 = 6e_1$ if $r = 1/2$.

We will need one more observation from our computations of $\widehat{\Gamma}$:

- For knots with ≤ 16 crossings and $g(K) = 1$, all curves other than γ_0 are simple figure eights, and γ_0 has one of the three forms: the horizontal curve, the curve which is the invariant of the right-handed trefoil (see Figure 3), or the mirror of the right-handed trefoil curve (which is the invariant of the left-handed trefoil).

We will denote the three possibilities for γ_0 above as $\gamma_0^{\mathrm{horiz}}$, γ_0^{RHT}, and γ_0^{LHT}. Note that they are distinguished by the value of τ, which is 0, 1, and -1, respectively.

9: 41, 44	10: 33, 136, 146	11a: 333
12a: 1144	13a: -	14a: 17464
15a: 76589, 84220	16a: 345268, 345454, 374264	11n: 18, 42, 62, 83

12n: 34, 65, 278, 313, 360, 393, 430, 483, 550, 650, 846, 884

13n: 71, 198, 490, 1019, 1209, 1398, 1513, 1598, 1756, 1757, 2337, 2703, 2796, 3290, 3416, 3783, 4591

14n: 372, 971, 1193, 2087, 2489, 6421, 7228, 7412, 7469, 7534, 8091, 8196, 8554, 8716, 9290, 9684, 9829, 10155, 11129, 11429, 12224, 12609, 12977, 13570, 14799, 15285, 15380, 15581, 15965, 15976, 17163, 17183, 18494, 19673, 21231, 21269, 22150, 22196, 22614, 22634, 23325, 24593, 27072, 27091

15n: 1058, 3240, 4898, 9477, 11491, 19192, 21666, 21997, 27824, 30711, 34041, 34773, 36113, 38567, 38594, 41604, 43982, 46350, 46536, 49081, 51379, 51847, 54458, 58840, 62260, 63468, 63550, 64468, 67694, 67879, 71170, 73390, 73507, 76978, 77245, 77247, 77784, 83761, 84434, 84645, 88899, 91448, 93899, 94474, 96914, 97157, 102309, 104775, 105829, 106611, 118711, 120250, 124511, 129229, 129231, 132539, 135706, 137623, 140373, 140582, 142082, 142299, 142716, 142841, 142843, 143482, 143825, 143856, 144436, 144439, 144887, 147186, 156806, 160027

16n: 5596, 9193, 16004, 24365, 27992,49009, 60136, 67523, 94939, 102539, 102773, 191694, 196472, 197735, 203049, 215168, 218032, 219174, 220556, 227624, 230857, 233335, 239267, 239379, 242042, 242545, 249927, 265957, 271606, 271610, 273164, 277974, 280482, 285128, 306917, 307635, 315594, 324571, 329529, 332372, 349983, 353272, 360174, 366612, 376208, 385669, 386732, 387806, 401152, 401963, 402644, 405088, 412371, 423420, 424451, 429723, 438719, 440479, 441595, 459035, 460502, 461585, 463225, 463419, 465019, 466470, 467558, 469510, 470606, 470717, 473737, 475444, 481843, 489486, 493489, 494163, 498542, 498651, 508893, 513585, 515663, 534392, 540621, 544661, 550305, 551107, 577882, 585135, 587843, 588588, 596192, 596449, 597513, 598535, 599034, 600112, 606009, 608181, 609311, 609798, 614804, 614907, 617672, 628265, 629526, 631987, 632225, 635338, 666646, 687419, 691300, 696924, 696992, 725574, 761555, 762559, 767010, 768788, 770126, 774829, 784110, 786382, 788898, 789181, 798964, 809799, 810368, 812243, 824554, 828723, 847911, 855704, 855909, 862009, 863179, 864017, 864258, 864259, 869439, 869441, 874997, 879694, 880152, 888060, 888954, 902353, 906603, 907441, 907673, 916183, 916207, 916242, 918157, 919068, 925408, 932460, 941562, 941564, 968742, 972142, 988939, 989795, 990225, 990270, 991069, 991085, 998071, 1000650, 1000651, 1001406, 1001474, 1004278, 1004646

Tab. 1. Knots up to 16 crossing for which Heegaard Floer homology does not rule out all cosmetic surgeries. The only pairs of slopes not ruled out for each of these knots are ± 1 and ± 2.

If $n = 2$, we must have $K = K_1 \# K_2$ with $g(K_1) = g(K_2) = 1$. By the Künneth formula for knot Floer homology, $\widehat{\Gamma}(K) = \widehat{\Gamma}(K_1) \otimes \widehat{\Gamma}(K_2)$[4]. A straightforward computation reveals that the tensor product of two figure eight components (at height 0) yields four figure eight components at heights $1, 0, 0$, and -1. The tensor product of a figure eight component and either γ_0^{RHT} or γ_0^{LHT} gives three figure eight components at heights $-1, 0$, and 1, and the tensor product of γ_0^{RHT} with γ_0^{LHT} yields a horizontal curve γ_0 along with two figure eights at heights -1 and 1. Taking the tensor product of any curve with $\gamma_0^{\mathrm{horiz}}$ gives a copy of that curve. Let a and b be the number of height 1 figure eight components in $\widehat{\Gamma}(K_1)$ and $\widehat{\Gamma}(K_2)$. Since τ is additive and $\tau(K) = 0$, either the γ_0 curves for K_1 and K_2 are γ_0^{RHT} and γ_0^{LHT} (in that order, without loss of generality) or they are both horizontal and $a, b > 0$. In the first case, we see that $\widehat{\Gamma}(K)$ has $e_0 = 2ab + a + b$ and $e_1 = ab + a + b + 1$; this is impossible since then $e_0 < 2e_1$. In the second case we have $e_0 = 2ab + a + b$ and $e_1 = ab$; this is impossible because $\frac{e_0}{e_1} = 2 + \frac{1}{a} + \frac{1}{b}$ is strictly greater than 2 and strictly less than 6. Thus $n \neq 2$.

If $n = 3$, we have $K = K_1 \# K_2 \# K_3$ with $g(K_i) = 1$. Let a, b, and c be the numbers of figure eight components in $\widehat{\Gamma}$ for K_1, K_2, and K_3. Since $\tau(K) = 0$, either γ_0 for the three knots is given (up to reordering) by γ_0^{RHT}, γ_0^{LHT}, and $\gamma_0^{\mathrm{horiz}}$, or γ_0 is horizontal for all three knots. In the first case, we compute that $e_0 = 6abc + 4ac + 4bc + 3c + 2ab + a + b$, and $e_1 = 4abc + 3ac + 3bc + 2c + ab + a + b + 1$. It follows that $2e_1 - e_0 = 2abc + 2ac + 2bc + c + a + b + 2 > 0$, which is a contradiction because we require $e_0 = 2e_1 + 7e_2 \geq 2e_1$. In the second case, we compute that $e_0 = 4abc + 2(ab + ac + bc) + a + b + c$, $e_1 = 3abc + (ab + ac + bc)$, and $e_2 = abc$. It follows that $7e_2 + 2e_1 > e_0$, since $7e_2 + 2e_1 - e_0 = 9abc - (a + b + c)$ is strictly positive. Thus $n \neq 3$. ∎

Remark 40. Note that in the above proof, the thickness bound immediately ruled out connected sums of more than three 16 crossing knots. Since $\mathrm{th}(K)$ and $g(K)$ are both additive and the upper bound on q goes like th/g^2, this behavior is expected. In fact, for any finite set of knots Theorem 2 prohibits truly cosmetic surgeries on connected sums of sufficiently many knots in the initial set.

Acknowledgments. I am grateful to Dave Futer for sparking my interest in this problem and for helpful correspondence, and to Liam Watson for comments on an earlier draft of this paper.

Funding. The author was partially supported by NSF grant DMS-1812527.

References

[1] Boyer, S., Lines, D.: Surgery formulae for Casson's invariant and extensions to homology lens spaces. J. Reine Angew. Math. **405**, 181–220 (1990) Zbl 0691.57004 MR 1041002

[4] This is an abuse of notation, as we really mean the tensor product of the corresponding bifiltered chain complexes; an immersed curve description of the tensor product operation is given in [9, Section 4]

[2] Futer, D., Purcell, J. S., Schleimer, S.: Effective bilipschitz bounds on drilling and filling. Geom. Topol., to appear

[3] Gainullin, F.: The mapping cone formula in Heegaard Floer homology and Dehn surgery on knots in S^3. Algebr. Geom. Topol. **17**, 1917–1951 (2017) Zbl 1380.57014 MR 3685598

[4] Gordon, C. M.: Dehn surgery on knots. In: Proceedings of the International Congress of Mathematicians, Vol. I (Kyoto, 1990), Math. Soc. Japan, Tokyo, 631–642 (1991) Zbl 0743.57008 MR 1159250

[5] Gordon, C. M., Luecke, J.: Knots are determined by their complements. J. Amer. Math. Soc. **2**, 371–415 (1989) Zbl 0672.57009 MR 965210

[6] Hanselman, J.: Bordered Heegaard Floer homology and graph manifolds. Algebr. Geom. Topol. **16**, 3103–3166 (2016) Zbl 1364.57012 MR 3584255

[7] Hanselman, J.: Knot floer homology via immersed curves. In preparation

[8] Hanselman, J., Rasmussen, J., Watson, L.: Bordered Floer homology for manifolds with torus boundary via immersed curves. arXiv:1604.03466 (2016)

[9] Hanselman, J., Rasmussen, J., Watson, L.: Heegaard Floer homology for manifolds with torus boundary: properties and examples. arXiv:1810.10355 (2018)

[10] Hanselman, J., Watson, L.: A calculus for bordered Floer homology. Geom. Topol., to appear

[11] Hom, J.: Bordered Heegaard Floer homology and the tau-invariant of cable knots. J. Topol. **7**, 287–326 (2014) Zbl 1368.57002 MR 3217622

[12] Hom, J.: A survey on Heegaard Floer homology and concordance. J. Knot Theory Ramifications **26**, art. 1740015, 24 pp. (2017) Zbl 1360.57002 MR 3604497

[13] Ichihara, K., Wu, Z.: A note on Jones polynomial and cosmetic surgery. Comm. Anal. Geom. **27**, 1087–1104 (2019) Zbl 1436.57005 MR 4037376

[14] Kirby, R.: Problems in low dimensional manifold theory. In: Algebraic and Geometric Topology (Stanford, CA, 1976), Part 2, Proc. Sympos. Pure Math. 32, Amer. Math. Soc., Providence, RI, 273–312 (1978) Zbl 0394.57002 MR 520548

[15] Ni, Y., Wu, Z.: Cosmetic surgeries on knots in S^3. J. Reine Angew. Math. **706**, 1–17 (2015) Zbl 1328.57010 MR 3393360

[16] Ozsváth, P., Szabó, Z.: Absolutely graded Floer homologies and intersection forms for four-manifolds with boundary. Adv. Math. **173**, 179–261 (2003) Zbl 1025.57016 MR 1957829

[17] Ozsváth, P., Szabó, Z.: Holomorphic disks and knot invariants. Adv. Math. **186**, 58–116 (2004) Zbl 1062.57019 MR 2065507

[18] Ozsváth, P., Szabó, Z.: On knot Floer homology and lens space surgeries. Topology **44**, 1281–1300 (2005) Zbl 1077.57012 MR 2168576

[19] Ozsváth, P. S., Szabó, Z.: Knot Floer homology and rational surgeries. Algebr. Geom. Topol. **11**, 1–68 (2011) Zbl 1226.57044 MR 2764036

[20] Rasmussen, J.: Lens space surgeries and a conjecture of Goda and Teragaito. Geom. Topol. **8**, 1013–1031 (2004) Zbl 1055.57010 MR 2087076

[21] Rasmussen, J. A.: Floer homology and knot complements. PhD thesis, Harvard Univ. (2003) MR 2704683

[22] Tao, R.: Cable knots do not admit cosmetic surgeries. J. Knot Theory Ramifications **28**, art. 1950034, 11 pp. (2019) Zbl 1429.57007 MR 3943698

[23] Wang, J.: Cosmetic surgeries on genus one knots. Algebr. Geom. Topol. **6**, 1491–1517 (2006) Zbl 1130.57020 MR 2253457

[24] Wu, Z.: Cosmetic surgery in L-space homology spheres. Geom. Topol. **15**, 1157–1168 (2011) Zbl 1226.57016 MR 2831258

[25] Xiu, Y.: Elliptic involution on knot complements. arXiv:1511.02298 (2015)